Compressive Classification of a Mixture of Gaussians: Analysis, Designs and Geometrical Interpretation

Hugo Reboredo, *Student Member, IEEE,* Francesco Renna, *Member, IEEE,* Robert Calderbank, *Fellow, IEEE,* and Miguel R. D. Rodrigues, *Member, IEEE*

arXiv:1401.6962v1 [cs.IT] 27 Jan 2014

Abstract

This paper derives fundamental limits on the performance of compressive classification when the source is a mixture of Gaussians. It provides an asymptotic analysis of a Bhattacharya based upper bound on the misclassification probability for the optimal Maximum-A-Posteriori (MAP) classifier that depends on quantities that are dual to the concepts of diversity-order and coding gain in multi-antenna communications. The diversity-order of the measurement system determines the rate at which the probability of misclassification decays with signal-to-noise ratio (SNR) in the low-noise regime. The counterpart of coding gain is the measurement gain which determines the power offset of the probability of misclassification in the low-noise regime. These two quantities make it possible to quantify differences in misclassification probability between random measurement and (diversity-order) optimized measurement. Results are presented for two-class classification problems first with zero-mean Gaussians then with nonzero-mean Gaussians, and finally for multiple-class Gaussian classification problems. The behavior of misclassification probability is revealed to be intimately related to certain fundamental geometric quantities determined by the measurement system, the source and their interplay. Numerical results, representative of compressive classification of a mixture of Gaussians, demonstrate alignment of the actual misclassification probability with the Bhattacharya based upper bound. The connection between the misclassification performance and the alignment between source and measurement geometry may be used to guide the design of dictionaries for compressive classification.

Index Terms

Compressed sensing, compressive classification, reconstruction, classification, random projections, measurement design, dimensionality reduction, Gaussian mixture models, phase transitions, diversity gain, measurement gain.

This paper was presented in part at the 2013 IEEE International Symposium on Information Theory and the 2013 IEEE Global Conference on Signal and Information Processing. The work of H. Reboredo was supported by Fundação para a Ciência e Tecnologia, Portugal, through the doctoral grant SFRH/BD/81543/2011. The work of M. R. D. Rodrigues was supported by the EPSRC through the research grant EP/K503459/1. This work was also supported by the Royal Society International Exchanges Scheme IE120996.

H. Reboredo and F. Renna are with the Instituto de Telecomunicações and the Departamento de Ciência de Computadores da Faculdade de Ciências da Universidade do Porto, Portugal (email: {hugoreboredo, frarenna}@dcc.fc.up.pt).

R. Calderbank is with the Department of Electrical and Computer Engineering, Duke University, NC, USA (email: robert.calderbank@duke.edu).

M. R. D. Rodrigues is with the Department of Electronic and Electrical Engineering, University College London, United Kingdom (email: m.rodrigues@ucl.ac.uk).

I. Introduction

A. The Compressive Classification Problem

Compressive sensing (CS) is an emerging paradigm that offers the means to simultaneously sense and compress a signal without any loss of information [1]–[4]. The sensing process is based on the projection of the signal of interest onto a set of vectors, which are typically constituted randomly [1]–[6], and the recovery process is based on the resolution of an inverse problem. The result that has captured the imagination of the signal and information processing community is that it is possible to perfectly reconstruct an n-dimensional s-sparse signal (sparse in some orthonormal dictionary or frame) with overwhelming probability with only $\mathcal{O}\left(s \log\left(n/s\right)\right)$ linear random measurements or projections [1], [4], [5] using tractable ℓ_1 minimization methods [2] or iterative methods, like greedy matching pursuit [7]–[9]. As such, compressive sensing has been proposed for a myriad of applications ranging from signal, image and video compression and processing, to communications and to medicine [10]–[13].

The focus of compressive sensing has been primarily on exact or near-exact signal reconstruction from a set of linear signal measurements. However, it is also natural to leverage the paradigm to perform other relevant information processing tasks, such as detection, classification and estimation of certain parameters, from the set of compressive measurements. One could in fact argue that the paradigm is a better fit to decision support tasks such as signal detection, signal classification or pattern recognition rather than signal reconstruction, since it may be easier to discriminate between signal classes than reconstruct an entire signal using only partial information about the source signal. Compressive information processing, recently proposed by Davenport *et al.* [14], thus advocates the resolution of various information processing tasks directly in the compressive measurement domain rather than the original possibly high-dimensional signal domain, which would entail resorting to full-scale signal reconstruction. In particular, Davenport *et al.* [14] argue that the use of the conventional random compressive measurement strategies that are agnostic to the exact form of the original signal – so applicable to a large class of signals – is key to develop very efficient and flexible compressive sensing hardware that can be used for the acquisition and processing of a large variety of signals.

This paper aims to study in detail the performance of a particular compressive information processing task: the classification of (possibly high-dimensional) signals from a set of compressive linear and noisy measurements. This problem is fundamental to the broad fields of signal and image processing [15]–[17], computer vision [18], [19] and machine learning [20]–[22], and pre-processing often relies on dimension reduction to increase the speed and reliability of classification as well as reduce the complexity and cost of data processing and computation. The question then becomes that of understanding how compressive measurements affect the classification performance as a function of the model parameters.

B. Prior Work on Compressive Classification

Compressive classification appears in the machine learning literature as feature extraction or supervised dimensionality reduction. For example, linear dimensionality reduction methods based on geometrical characterizations

of the source have been developed, with linear discriminant analysis (LDA) [23] and principal component analysis (PCA) [23] just depending on second order statistics. In particular, LDA, which is one of the most well-known supervised dimensionality reduction methods [24], addresses simultaneously the between-class scattering and the within-class scattering of the projected data. It has been proven that under mild conditions this method is Bayes optimal [25]. However, this method has two disadvantages: i) the dimensionality of the projected space can only be less than the number of data classes, which greatly restricts its applicability; and ii) LDA only uses first and second order statistics of the data, ignoring the higher-order information. Linear dimensionality reduction methods based on higher-order statistics of the data have therefore also been developed [20], [22], [26]–[34]. In particular, an information-theoretic supervised dimensionality reduction inspired approach, which uses the mutual information between the data class labels and the data projections [22], [26] or approximations of the mutual information such as quadratic mutual information (with quadratic Rényi entropy) [28], [33], [34] as a criterion to linearly reduce dimensionality, have been shown to lead to state-of-the-art classification results. The rationale for using mutual information relates to the fact that the Bayes classification error is bounded by the mutual information (based on a Shannon entropy measure) between the data labels and the data projections [31]. In addition, nonlinear (supervised) dimensionality reduction methods have become popular recently [35].

In turn, [21] and [36] consider compressed learning, i.e. learning directly in the compressed domain rather than in the original data domain. Of particular relevance, bearing witness to the value of the compressive classification paradigm, it is shown that the linear kernel support vector machine (SVM) classifier (trained and working) in the measurement domain, with high probability, has true accuracy close to the accuracy of the best linear threshold classifier (trained and working) in the data domain. It is also shown that for a family of well-known compressed sensing matrices, compressed learning is universal, in the sense that learning and classification in the measurement domain works provided that the data are sparse in some, even unknown, basis – therefore, compressed sensing matrices also provide the desired properties of good linear dimensionality reduction matrices. That is, compressed sensing can be used as an efficient transformation from the data domain to the measurement domain (when the data admits a sparse representation even in an unknown basis) that preserves the learnability (and separability) of the data while bypassing the computational cost and curse of dimensionality.

Compressive classification also appears in the compressive information processing literature in view of recent advances in compressive sensing [14], [20], [37]–[44]. References [37] and [38] explore random compressive measurements to perform manifold-based image classification. References [39], [40], [41] and [42] study the performance of compressive detection and compressive classification. Reference [14] considers various compressive information processing problems, including compressive detection and compressive classification. References [43] and [44] consider the problem of detection of spectral targets based on noisy incoherent projections. Reference [20] notes that a small number of random measurements captures sufficient information to allow robust face recognition. The common thread in this line of research relates to the demonstration that the detection and classification problems can be solved directly in the measurement domain, without requiring the transformation of the data from the compressive to the original data domain, i.e. without requiring the reconstruction of the data.

Other works associated with compressive classification that have arisen in the computational imaging literature, and developed under the rubric of task-specific sensing, include [15]–[17], [45]–[48]. In particular, task-specific sensing, which advocates that the sensing procedure has to be matched to the task-specific nature of the sensing application, has been shown to lead to substantial gains in performance over compressive sensing in applications such as localization [45], target detection [15], (face) recognition [16], [17], and reconstruction [46].

Another instance of compressive classification appears under the guise of sparse support recovery (also known as model selection) associated with compressive sensing problems [49]–[61]. These works provide necessary and sufficient (high-dimensional) scalings on the triplet (n, s, m), where n relates to the signal dimension, s relates to the signal sparsity and m is the number of linear measurements, for successful or partially successful support recovery under various metrics, various decoders (optimal and sub-optimal decoders), various measurement matrices, and various sparsity regimes (e.g. the linear sparsity regime and the sub-linear sparsity regime). References [53], [54] and [55] provide a more refined analysis associated with the tradeoff between the number of measurements and the number of detection errors in the high-dimensional setting, where the sparsity rate (i.e. the fraction of nonzero entries) and the per-sample signal-to-noise ratio (SNR) are finite constants, independent of the signal dimension.

C. A Characterization of the Incremental Value of Measurement

This paper provides a finer grained characterization of the performance of a compressive classification system. It goes beyond the phase-transition inspired high-dimensional characterizations that describe as a function of the source signal and system parameters whether or not one can reliably recover or partially recover the sparse signal support (i.e. presence or absence of a misclassification probability floor) [49]–[61]. It quantifies the information that can be learned from incremental measurement as a function of SNR through the concepts of diversity-order and measurement gain.

a) The diversity-order determines the decay (in a $\log(\mathrm{SNR})$ scale) of the misclassification probability at low-noise levels (high-SNR levels);
b) The measurement gain determines the horizontal offset (in a $\log(\mathrm{SNR})$ scale) of the misclassification probability at low-noise levels (high-SNR levels), i.e. the measurement gain distinguishes further characterizations that exhibit identical diversity-order.

These metrics, which determine the power of measurement to discriminate classes, will be naturally described in terms of the geometry of the source and the measurement system, and this description will reveal how measurement leads to separation of classes. We suggest that these metrics be used as proxies when designing classification systems.

Enter wireless communications. The performance characterizations we propose for the compressive classification problem are parallel to those adopted in the wireless communications field [62]–[65]: diversity-order and coding gain in wireless communications problems are the counterparts of the diversity-order and measurement gain in the compressive classification problem. The reason for this correspondence is the signal model; the assumption that measurement noise is Gaussian and that the distribution of the source under each class (hypothesis) is multivariate Gaussian with a certain known mean and a certain known (possibly rank-deficient) covariance matrix. It is the

source model that leads to a fundamental duality between compressive classification and non-coherent wireless communication with multiple antennas [66]-[70]. This duality means that fundamental limits derived in the wireless domain can be transferred to the classification domain and vice versa.[1] There are two other notable features associated with this model:

a) The source model adopted for the compressive classification problem also relates to well-known models adopted in compressive sensing, most notably the Gaussian mixture model (GMM). The GMM [48], [68]–[70], which is typically used in conjunction with the Bayesian compressed sensing formalism [71], relates to various well-known structured models in the literature including union of sub-spaces [72]–[75], wavelet trees [72], [76] or manifolds [68], [77], that aim to capture additional signal structure beyond primitive sparsity in order to yield further gains. For example, a (low-rank) GMM can be seen as a Bayesian counterpart of the union of subspaces model. In fact, a signal drawn from a (low-rank) GMM lies in a union of subspaces, where each subspace corresponds to the image of each class conditioned covariance matrix in the model.[2] A low-rank GMM can also be seen as an approximation to a compact manifold. Compact manifolds can be covered by a finite collection of topological disks that can be represented by high-probability ellipsoids living on the principal hyperplanes corresponding to the different components of a low-rank GMM [68]. However, one of the key advantages of adopting a GMM, *in lieu* of these other structured models, is associated with the fact that reconstruction of a signal drawn from a GMM from compressive linear measurements in Gaussian noise can be very effectively performed via a closed-form inversion formula [68].
b) This source model also leads to state-of-the-art results in various compressive classification scenarios such as character and digit recognition as well as image classification [22], [33], [34]. Other successful instances – beyond classification – associated with the use of GMMs include problems in image processing, such as interpolation, zooming, deblurring [48], [69], [70], and problems in dictionary learning [68].

D. Contributions of the Article

The main contributions include:

1) Characterization of the behavior of an upper bound to the probability of misclassification for linear random Gaussian measurements contaminated by white Gaussian noise. The characterization unveils how the error floor, the diversity-order and the measurement gain behave as a function of the properties of the source (i.e. the number of classes and the means and covariances of the class conditioned multivariate Gaussian distributions);

[1]These dualities also lead to characterizations that disclose fundamental tradeoffs between the performance of a compressive classification problem (via the diversity-order) and the number of classes in the compressive classification problem (a so-called discrimination gain). Such a diversity-discrimination tradeoff [66], which offers a more illuminating view of the performance of the compressive classification, is reminiscent of the diversity-multiplexing tradeoff in multiple-antenna wireless communications [65], [67].

[2]More generally, a signal drawn from a GMM model lies in a union of affine spaces rather than linear subspaces, where each affine space is associated with the mean and covariance of each class in the GMM model.

2) Characterization of the behavior of an upper bound to the probability of misclassification for linear designed measurements contaminated by Gaussian noise. In particular, we construct projection designs that maximize the diversity-order subject to a certain measurement budget as a function of the properties of the source;
3) Extension of the performance characterizations from the low-noise to the high-noise regime. Such an extension showcases the key operational differences between the two regimes;
4) Connection of the performance behavior of the compressive classification problem, with random or designed measurements, to certain geometrical quantities associated with the measurement system and associated with the source, in order to provide additional insight.

These contributions differ from other related contributions in the prior work on compressive classification (including the literature on sparse support recovery) in various aspects:

1) Prior work on performance measures has focused on memoryless source models (e.g. the Bernoulli-Gaussian model or the models akin to the Bernoulli-Gaussian model associated with sparse support recovery problems) [42], [49]–[61], [78]; in contrast, we define the behavior of the probability of misclassification for source models that inherently incorporate memory;
2) Prior analysis of performance measures is typically asymptotic in various of the problem dimensions (e.g. the signal dimension, the signal sparsity, the number of measurements and certain signal parameters for certain scalings) [49]–[61], [78]; in contrast, we define the behavior of the probability of misclassification non-asymptotically in various of the problem dimensions and asymptotically only in the signal-to-noise ratio;
3) Prior work has concentrated on identifying phase transitions, whereas we concentrate on characterizations that articulate about the presence of absence of an error floor, the diversity-order and the measurement gain.

Some elements of our approach are present in the characterizations presented in [14], [40]–[42] but their focus is not on the incremental value of measurement to classification. In particular, the additional refinement offered by our approach can also be appreciated by specializing our setting, where the source signal lives in a union of sub-spaces, to the traditional sparse recovery setting, where the source signals are taken to be sparse in the canonical basis. For example, let us consider a n-dimensional GMM source formed by $\sum_{i=1}^{k} \binom{n}{i}$ equiprobable classes, where each Gaussian class is represented by a zero-mean vector and a diagonal covariance matrix with rank less than or equal to k, whose nonzero elements correspond to one out of the $\sum_{i=1}^{k} \binom{n}{i}$ possible supports of cardinality less than or equal to k in $\mathbb{R}^n$. The results in [58], [60] show that it is possible to recover the sparse signal support pattern (classify the signal sub-space) with $\mathcal{O}(k \log n)$ measurements[3]. On the other hand, our results show that it is possible to classify the signal sub-space with exactly k random measurements when the noise power tends to zero. In addition, our results also lead to a sharper characterization of the behavior of the misclassification probability: it decays with at least a slope of $1/4$ on a log(SNR) scale.

[3]Note that these results hold in particular when the minimum amplitude of the nonzero entries in the input vector is assumed to decrease as $\mathcal{O}(1/\sqrt{k})$ and for a fixed value of the noise power σ^2.

E. Structure of the Article

This article is structured as follows: Section II introduces the compressive classification problem and the associated system, signal and noise models. Sections III and IV study the performance of compressive classification for the two-class and the multiple-class scenarios, respectively, with random measurements in the regime of low noise. In turn, Section V studies the performance of two- and multiple-class compressive classification problems with designed measurements also in the low-noise regime. The purpose of Section VI is to extend the analysis from the low-noise to the high-noise regime. A comprehensive set of numerical results that sheds further light on the performance of compressive classification problems is provided in Section VII. Concluding remarks are made in Section VIII. The technical proofs of the main results are delegated to the Appendices.

F. Notation

The article adopts the following notation: boldface upper-case letters denote matrices ($\mathbf{X}$), boldface lower-case letters denote column vectors ($\mathbf{x}$) and italics denote scalars (x); the context defines whether the quantities are deterministic or random. $\mathbf{I}_N$ represents the $N \times N$ identity matrix, $\mathbf{0}_{M\times N}$ represents the $M \times N$ zero matrix (the subscripts that refer to the dimensions of such matrices will be droppep when evident from the context) and $\mathrm{diag}\left(a_1, a_2, \ldots, a_N\right)$ represents an $N \times N$ diagonal matrix with diagonal elements $a_1, a_2, \ldots, a_N$. The operators $\left(\cdot\right)^T$, $\mathrm{rank}\left(\cdot\right)$, $\det\left(\cdot\right)$, $\mathrm{pdet}\left(\cdot\right)$ and $\mathrm{tr}\left(\cdot\right)$ represent the transpose operator, the rank operator, the determinant operator, the pseudo-determinant operator and the trace operator, respectively. $\mathrm{Null}\left(\cdot\right)$ and $\mathrm{im}\left(\cdot\right)$ denote the null space and the (column) image of a matrix, respectively, and $\dim\left(\cdot\right)$ denotes the dimension of a linear subspace. $\mathcal{E}\left\{\cdot\right\}$ represents the expectation operator. The multi-variate Gaussian distribution with mean $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$ is denoted by $\mathcal{N}\left(\boldsymbol{\mu}, \boldsymbol{\Sigma}\right)$. $\log\left(\cdot\right)$ denotes the natural logarithm. The article also uses the big $\mathcal{O}$ notation where $g\left(x\right) = \mathcal{O}\left(f\left(x\right)\right)$ if $\lim_{x\to\infty} \frac{g\left(x\right)}{f\left(x\right)} = c$, where c is a constant; and the little o notation where $g\left(x\right) = o\left(f\left(x\right)\right)$ if $\lim_{x\to\infty} \frac{g\left(x\right)}{f\left(x\right)} = 0$.

II. The Compressive Classification Problem

We consider the standard measurement model given by:

$$\mathbf{y} = \boldsymbol{\Phi}\mathbf{x} + \mathbf{n}, \tag{1}$$

where $\mathbf{y} \in \mathbb{R}^M$ represents the measurement vector, $\mathbf{x} \in \mathbb{R}^N$ represents the source vector, $\boldsymbol{\Phi} \in \mathbb{R}^{M\times N}$ represents the measurement matrix or kernel[4] and $\mathbf{n} \sim \mathcal{N}\left(\mathbf{0}, \sigma^2 \cdot \mathbf{I}\right) \in \mathbb{R}^M$ represents standard white Gaussian noise. We consider both random measurement kernel designs, where the entries of $\boldsymbol{\Phi}$ are drawn i.i.d. from a zero-mean, fixed-variance, Gaussian distribution, which is common in the CS literature [1], [4] as well as optimized kernel designs that aim to improve the classification performance.

We suppose that the source is described by a Gaussian Mixture Model (GMM); the source signal is drawn from one of L classes, the *a priori* probability of class i is P_i and the distribution of the source conditioned on class

[4]We refer to $\boldsymbol{\Phi}$ as the measurement, sensing or projection matrix/kernel interchangeably throughout the paper.

i is Gaussian with mean $\boldsymbol{\mu}_i \in \mathbb{R}^N$ and (possibly rank-deficient) covariance matrix $\boldsymbol{\Sigma}_i \in \mathbb{R}^{N\times N}$. We should point out that we use a low-rank modelling approach even though natural signals (e.g. patches extracted from natural images) are not always low-rank but rather "approximately" low-rank [68]. The justification for the use of such low-rank modeling approach is two-fold: first, a low-rank representation is often a very good approximation to real scenarios, particularly as the eigenvalues of the class conditioned covariances often decay rapidly; second, it is then standard practice to account for the mismatch between the low-rank and the "approximately" low-rank model by adding extra noise in the measurement model in (1) (see [79]).

Our objective is then to produce an estimate of the true signal class given the measurement vector. The Maximum-A-Posteriori (MAP) classifier, which minimizes the probability of misclassification [23], produces the estimate given by:

$$\hat{C} = \arg \max_{C\in\{1,\cdots,L\}} P\left(C \mid \mathbf{y}\right) = \arg \max_{C\in\{1,\cdots,L\}} p\left(\mathbf{y} \mid C\right) P_C, \tag{2}$$

where $P(C \mid \mathbf{y})$ represents the *a posteriori* probability of class C given the measurement vector $\mathbf{y}$ and $p(\mathbf{y} \mid C)$ represents the probability density function of the measurement vector $\mathbf{y}$ given the class C.

Our performance analysis concentrates both on the low-noise regime, where $\sigma^2 \to 0$, and on the high-noise regime, where $\sigma^2 \to \infty$. The performance analysis – in line with the standard practice in multiple-antenna communications systems [62], [64] – also concentrates on the asymptotic behavior of an upper bound to the probability of misclassification of the MAP classifier P_{err}^{UB} , rather than the exact probability of misclassification of such classifier P_{err}. In the high-noise regime, we use standard Taylor series expansions to describe the asymptotic behavior of the probability of misclassification.

In contrast, in the low-noise regime, we use two measures that describe the low-noise asymptotics of the upper bound to the probability of misclassification. In particular, we define the (lower bound to the) diversity-order of the measurement model in (1) as:

$$d = \lim_{\sigma^2\to 0} \frac{\log P_{err}^{UB}(\sigma^2)}{\log \sigma^2}, \tag{3}$$

that determines how (the upper bound to) the misclassification probability decays (in the $\log \sigma^2$ scale) at low-noise levels [65], [80]. We also define the (lower bound to the) measurement gain of the measurement model in (1) as:

$$g_m = \lim_{\sigma^2\to 0} \sigma^2 \cdot \frac{1}{\sqrt[d]{P_{err}^{UB}(\sigma^2)}}, \tag{4}$$

that determines the power offset of (the upper bound to) the misclassification error probability at low-noise levels: note that the measurement gain refines the asymptotic description of the upper bound to the misclassification probability, by distinguishing further characterizations that exhibit identical diversity-order.

The characterization of the performance measures in (3) and (4) will be expressed via quantities that relate to the geometry of the measurement model, namely, the rank and the pseudo-determinant of certain matrices. In particular, we define the behavior of (3) and (4) in two ways. The geometric interplay of the linear transformation with the class covariance matrices is described by the quantities::

$$r_i = \operatorname{rank}(\boldsymbol{\Phi}\boldsymbol{\Sigma}_i\boldsymbol{\Phi}^T) \tag{5}$$

and

$$v_i = \mathrm{pdet}(\mathbf{\Phi}\mathbf{\Sigma}_i\mathbf{\Phi}^T) \tag{6}$$

which measure the dimension of the sub-space spanned by the linear transformation of the signals in class i and the volume of the parallelepiped in $\mathbb{R}^M$ where those signals are mostly concentrated, respectively; and the quantities:

$$r_{ij} = \mathrm{rank}(\mathbf{\Phi}(\mathbf{\Sigma}_i + \mathbf{\Sigma}_j)\mathbf{\Phi}^T) \tag{7}$$

and

$$v_{ij} = \mathrm{pdet}(\mathbf{\Phi}(\mathbf{\Sigma}_i + \mathbf{\Sigma}_j)\mathbf{\Phi}^T) \tag{8}$$

which measure the dimension of the direct sum of sub-spaces spanned by the linear transformation of the signals in classes i or j and the volume of the parallelepiped in $\mathbb{R}^M$ where signals in classes i and j are mostly concentrated, respectively.

In addition, we also define the behavior of the measures in (3) and (4) via the geometry of the original source signal, by using the quantities:

$$r_{\mathbf{\Sigma}_i} = \mathrm{rank}(\mathbf{\Sigma}_i) \tag{9}$$

which relates to the dimension of the sub-space spanned by source signals in class i and

$$r_{\mathbf{\Sigma}_{ij}} = \mathrm{rank}(\mathbf{\Sigma}_i + \mathbf{\Sigma}_j) \tag{10}$$

which relates to the dimension of the direct sum of sub-spaces spanned by source signals in classes i or j. Finally, we will also be using the quantity:

$$NO_{Dim} = r_{\mathbf{\Sigma}_{ij}} - \left[\left(r_{\mathbf{\Sigma}_i} + r_{\mathbf{\Sigma}_j}\right) - r_{\mathbf{\Sigma}_{ij}}\right] \tag{11}$$

that relates to the difference between the dimension of the sub-spaces spanned by source signals in classes i or j and the dimension of the intersection of such sub-spaces. This can also be interpreted as the number of non-overlapping dimensions between the sub-spaces spanned by the eigenvectors of the covariance matrices pertaining to the two classes.

It turns out - as shown in the sequel - that the asymptotics of the upper bound to the misclassification probability mimic closely the behavior of the true misclassification probability, attesting to the value of the approach.

III. Two-Class Compressive Classification with Random Measurements in the Low-Noise Regime

We set $L = 2$ and consider two-class compressive classification using random measurements. The misclassification probability associated to the MAP classifier can be written as

$$P_{err} = \int_{-\infty}^{+\infty} \min\left(P_1 \cdot p(\mathbf{y}|C=1), P_2 \cdot p(\mathbf{y}|C=2)\right) d\mathbf{y}. \tag{12}$$

and, by exploiting the fact that $\min(a,b) \leq a^t b^{1-t}, \;\; 0 \leq t \leq 1, \;\; a,b>0$, an upper bound to the misclassification probability can also be immediately written as [23]:

$$P_{err} \leq P_1^t P_2^{1-t} \int_{-\infty}^{+\infty} p^t(\mathbf{y} \mid C=1)\, p^{1-t}(\mathbf{y} \mid C=2)\, d\mathbf{y}, \quad 0 \leq t \leq 1. \tag{13}$$

It turns out that the Bhattacharyya bound [81], which corresponds to a specialization of the bound in (13) for $t=0.5$, admits the closed-form expression for class-conditioned Gaussian distributions given by: [5]

$$P_{err}^{UB} = \sqrt{P_1 P_2} \int_{-\infty}^{+\infty} \sqrt{p(\mathbf{y} \mid C=1)\, p(\mathbf{y} \mid C=2)} d\mathbf{y} = \sqrt{P_1 P_2}\; e^{-K_{12}}, \tag{14}$$

where

$$\begin{aligned} K_{ij} = \frac{1}{8}\left[\mathbf{\Phi}\left(\boldsymbol{\mu}_i - \boldsymbol{\mu}_j\right)\right]^T \left[\frac{\mathbf{\Phi}\left(\mathbf{\Sigma}_i + \mathbf{\Sigma}_j\right)\mathbf{\Phi}^T + 2\sigma^2\mathbf{I}}{2}\right]^{-1} \left[\mathbf{\Phi}\left(\boldsymbol{\mu}_i - \boldsymbol{\mu}_j\right)\right] \\ + \frac{1}{2}\log \frac{\det\left(\frac{\mathbf{\Phi}(\mathbf{\Sigma}_i+\mathbf{\Sigma}_j)\mathbf{\Phi}^T+2\sigma^2\mathbf{I}}{2}\right)}{\sqrt{\det\left(\mathbf{\Phi}\mathbf{\Sigma}_i\mathbf{\Phi}^T+\sigma^2\mathbf{I}\right)\det\left(\mathbf{\Phi}\mathbf{\Sigma}_j\mathbf{\Phi}^T+\sigma^2\mathbf{I}\right)}}. \end{aligned} \tag{15}$$

This Bhattacharyya based upper bound on the probability of misclassification encapsulated in (14) and (15) is the basis of our analysis. We treat the case where the classes are zero-mean, i.e. $\boldsymbol{\mu}_1 = \boldsymbol{\mu}_2 = \mathbf{0}$, and the case where classes are nonzero-mean, i.e. $\boldsymbol{\mu}_1 \neq \mathbf{0}$ or $\boldsymbol{\mu}_2 \neq \mathbf{0}$, separately. The zero-mean case exhibits the main operational features associated with the compressive classification problem; the nonzero-mean case can also exhibit additional operational features, e.g. infinite diversity-order.

A. Zero-Mean Classes

The following Theorem offers a view of the asymptotic behavior of the (upper bound to the) probability of misclassification for the two-class compressive classification problem with zero-mean classes, by leveraging directly the geometry of the linear transformation of the source signal effected by the measurement kernel. Note that, since the matrices $\mathbf{\Phi}\mathbf{\Sigma}_1\mathbf{\Phi}^T$ and $\mathbf{\Phi}\mathbf{\Sigma}_2\mathbf{\Phi}^T$ are positive semidefinite, it is straightforward to show that $\frac{r_1+r_2}{2} \leq r_{12}$; in addition, it is also possible to show that $\frac{r_1+r_2}{2} = r_{12}$ if and only if $\mathrm{im}(\mathbf{\Phi}\mathbf{\Sigma}_1\mathbf{\Phi}^T) = \mathrm{im}(\mathbf{\Phi}\mathbf{\Sigma}_2\mathbf{\Phi}^T)$, so that the two sub-spaces associated with the two classes overlap completely [79, Lemma 2].

Theorem 1: Consider the measurement model in (1) where $\mathbf{x} \sim \mathcal{N}(\mathbf{0}, \mathbf{\Sigma}_1)$ with probability P_1 and $\mathbf{x} \sim \mathcal{N}(\mathbf{0}, \mathbf{\Sigma}_2)$ with probability $P_2 = 1 - P_1$. Then, in the regime of low noise where $\sigma^2 \to 0$, the upper bound to the probability of misclassification in (14) behaves as:

- If $\frac{r_1+r_2}{2} = r_{12}$ then, $P_{err}^{UB} = \mathcal{O}(1), \quad \sigma^2 \to 0$;
- If $\frac{r_1+r_2}{2} < r_{12}$ then, $P_{err}^{UB} = \left(\frac{g_m}{\sigma^2}\right)^{-d} + o\left(\left(\frac{1}{\sigma^2}\right)^{-d}\right), \quad \sigma^2 \to 0$, where:

[5]Note that Bhattacharyya upper bound corresponds to a value of $t=0.5$ in (13) whereas the Chernoff bound corresponds to the value of t that minimizes (13).

$$d = -\frac{1}{2}\left(\frac{r_1 + r_2}{2} - r_{12}\right) \tag{16}$$

and

$$g_m = \left[2^{\frac{r_{12}}{2}}\sqrt{P_1 P_2}\left[\frac{v_{12}}{\sqrt{v_1 v_2}}\right]^{-\frac{1}{2}}\right]^{-\frac{1}{d}}. \tag{17}$$

Proof: The proof is presented in Appendix A. ■

It is now instructive to probe further onto the characterizations embodied in Theorem 1 to infer some operational features associated with the two-class compressive classification problem. Such characterization admits a very simple interpretation:

- If $\frac{r_1+r_2}{2} = r_{12}$, then the sub-spaces spanned by the linear transformation of the signals in classes $C = 1$ and $C = 2$ overlap completely – the upper bound to the misclassification probability exhibits an error floor because it is not possible to distinguish the classes perfectly as the noise level approaches zero;
- If $\frac{r_1+r_2}{2} < r_{12}$, then the sub-spaces spanned by the linear transformation of the signals in classes $C = 1$ and $C = 2$ do not overlap completely – the upper bound to the misclassification error probability (and the true error probability) does not exhibit an error floor as it is possible to distinguish the classes perfectly as the noise level approaches zero. The lower the degree of overlap, the higher the diversity-order – this is measured via the interplay of the various ranks, r_1, r_2 and r_{12} in terms of the difference between the dimensions of the sub-spaces corresponding to the linear transformation of the signals in classes $C = 1$ and $C = 2$ and the dimension of their intersection; in fact, it can be shown that:

$$\begin{aligned} 2r_{12} - r_1 - r_2 = &\dim \operatorname{im}\left(\mathbf{\Phi}\mathbf{\Sigma}_1\mathbf{\Phi}^T\right) - \dim\left(\operatorname{im}\left(\mathbf{\Phi}\mathbf{\Sigma}_1\mathbf{\Phi}^T\right) \cap \operatorname{im}\left(\mathbf{\Phi}\mathbf{\Sigma}_2\mathbf{\Phi}^T\right)\right) \\ &+ \dim \operatorname{im}\left(\mathbf{\Phi}\mathbf{\Sigma}_2\mathbf{\Phi}^T\right) - \dim\left(\operatorname{im}\left(\mathbf{\Phi}\mathbf{\Sigma}_1\mathbf{\Phi}^T\right) \cap \operatorname{im}\left(\mathbf{\Phi}\mathbf{\Sigma}_2\mathbf{\Phi}^T\right)\right). \end{aligned} \tag{18}$$

The following Theorem now describes the asymptotic behavior of the probability of misclassification for the two-class compressive classification problem with zero-mean classes, by leveraging directly the geometry of the source signals – this has the advantage of showcasing how the number of measurements together with the properties of the source affect performance. The result uses the fact that $N \geq r_{\mathbf{\Sigma}_{12}} \geq \max\left(r_{\mathbf{\Sigma}_1}, r_{\mathbf{\Sigma}_2}\right)$ and, with probability 1, $r_1 = \min\left(M, r_{\mathbf{\Sigma}_1}\right)$, $r_2 = \min\left(M, r_{\mathbf{\Sigma}_2}\right)$ and $r_{12} = \min\left(M, r_{\mathbf{\Sigma}_{12}}\right)$. The result also assumes, without loss of generality, that $r_{\mathbf{\Sigma}_1} \leq r_{\mathbf{\Sigma}_2}$. Once more, note that since the matrices $\mathbf{\Sigma}_1$ and $\mathbf{\Sigma}_2$ are positive semidefinite, it follows that $\frac{r_{\mathbf{\Sigma}_1}+r_{\mathbf{\Sigma}_2}}{2} \leq r_{\mathbf{\Sigma}_{12}}$, and that $\frac{r_{\mathbf{\Sigma}_1}+r_{\mathbf{\Sigma}_2}}{2} = r_{\mathbf{\Sigma}_{12}}$ if and only if $\operatorname{im}(\mathbf{\Sigma}_1) = \operatorname{im}(\mathbf{\Sigma}_2)$ [79, Lemma 2].

Theorem 2: Consider the measurement model in (1) where $\mathbf{x} \sim \mathcal{N}(\mathbf{0}, \mathbf{\Sigma}_1)$ with probability P_1 and $\mathbf{x} \sim \mathcal{N}(\mathbf{0}, \mathbf{\Sigma}_2)$ with probability $P_2 = 1 - P_1$. Then, in the regime of low noise where $\sigma^2 \to 0$, the upper bound to the probability of misclassification in (14) behaves as:

- If $\frac{r_{\mathbf{\Sigma}_1}+r_{\mathbf{\Sigma}_2}}{2} = r_{\mathbf{\Sigma}_{12}}$ then, $P_{err}^{UB} = \mathcal{O}(1), \quad \sigma^2 \to 0$;

- If $\frac{r_{\mathbf{\Sigma}_1}+r_{\mathbf{\Sigma}_2}}{2} < r_{\mathbf{\Sigma}_{12}}$ then,
  - when $M \leq r_{\mathbf{\Sigma}_1} \leq r_{\mathbf{\Sigma}_2} \leq r_{\mathbf{\Sigma}_{12}}$, $P_{err}^{UB} = \mathcal{O}\left(1\right)$, $\sigma^2 \to 0$;
  - otherwise, $P_{err}^{UB} = \left(\frac{g_m}{\sigma^2}\right)^{-d} + o\left(\left(\frac{1}{\sigma^2}\right)^{-d}\right)$, $\sigma^2 \to 0$ where the measurement gain is given by:

$$g_m = \left[2^{\frac{\min(M, r_{\mathbf{\Sigma}_{12}})}{2}} \sqrt{P_1 P_2} \left[\frac{v_{12}}{\sqrt{v_1 v_2}}\right]^{-\frac{1}{2}}\right]^{-\frac{1}{d}} \tag{19}$$

  and the diversity-order is given by:

$$d = \begin{cases} -\frac{1}{2}\left(\frac{r_{\mathbf{\Sigma}_1} - M}{2}\right) & , \text{ if } r_{\mathbf{\Sigma}_1} < M \leq r_{\mathbf{\Sigma}_2} \leq r_{\mathbf{\Sigma}_{12}} \\ -\frac{1}{2}\left(\frac{r_{\mathbf{\Sigma}_1} + r_{\mathbf{\Sigma}_2}}{2} - M\right) & , \text{ if } r_{\mathbf{\Sigma}_1} \leq r_{\mathbf{\Sigma}_2} < M < r_{\mathbf{\Sigma}_{12}} \\ -\frac{1}{2}\left(\frac{r_{\mathbf{\Sigma}_1} + r_{\mathbf{\Sigma}_2}}{2} - r_{\mathbf{\Sigma}_{12}}\right) & , \text{ if } r_{\mathbf{\Sigma}_1} \leq r_{\mathbf{\Sigma}_2} \leq r_{\mathbf{\Sigma}_{12}} \leq M \end{cases} . \tag{20}$$

  Proof: The proof is presented in Appendix B. ■

Theorem 2 provides insight into the interplay between the number of measurements and the source geometry. In particular:

- When the sub-spaces spanned by the signals in classes $C = 1$ and $C = 2$ overlap completely, i.e., $\frac{r_{\mathbf{\Sigma}_1}+r_{\mathbf{\Sigma}_2}}{2} = r_{\mathbf{\Sigma}_{12}}$, it is not possible to construct a random measurement kernel that will be able to distinguish between the signals from classes $C = 1$ and $C = 2$. In such scenario, the (upper bound to the) probability of misclassification will exhibit an error floor at low-noise levels.
- When the sub-spaces spanned by the signals in classes $C = 1$ and $C = 2$ do not overlap completely, i.e., $\frac{r_{\mathbf{\Sigma}_1}+r_{\mathbf{\Sigma}_2}}{2} < r_{\mathbf{\Sigma}_{12}}$, the number of random measurements M defines the behavior of the (upper bound to the) probability of misclassification as follows:

  i) if $M \leq r_{\mathbf{\Sigma}_1}$, the upper bound will exhibit an error floor at low-noise levels, because the random measurements will affect the signals in a way that the sub-spaces associated with the linearly transformed classes completely overlap – that is, the random measurements do not provide the sufficient degrees of freedom to distinguish between the signals from classes $C = 1$ and $C = 2$. Note that in this case both the sub-spaces corresponding to the randomly projected classes occupy the entire space $\mathbb{R}^M$ thus they are completely overlapping;

  ii) otherwise, if $M > r_{\mathbf{\Sigma}_1}$ the upper bound will not exhibit such an error floor at low-noise levels, since it is possible to randomly transform the original signals such that the corresponding sub-spaces do not overlap completely; in fact, it can be shown directly from (18) that the difference between the dimensions of these sub-spaces and that of their intersection is at least one;

  iii) once again, note that the diversity-order is a function of the difference between the dimensions of the sub-spaces associated with the two classes and the dimension of their intersection, as given by (16) and (18): by gradually increasing the number of measurements M from 1 up to $r_{\mathbf{\Sigma}_{12}}$ it is possible to increase the diversity-order up to the maximum value $\frac{1}{4}NO_{Dim}$; however, increasing the number of measurements past $r_{\mathbf{\Sigma}_{12}}$ does not offer a higher diversity-order – instead, it only translates into a higher measurement gain. In

fact, as M ranges from 1 to $r_{\mathbf{\Sigma}_{12}}$ we can gradually unveil the "degree of non-overlap" between the original sub-spaces because the number of non-overlapping dimensions in the projected domain approaches the number of non-overlapping dimensions in the original domain, achieving it when $M = r_{\mathbf{\Sigma}_{12}}$. In other terms, when $M = r_{\mathbf{\Sigma}_{12}}$, the performance of classification – defined via the diversity-order – in the measurement domain equals that in the data domain [21], [36]. In contrast, for $M > r_{\mathbf{\Sigma}_{12}}$ projecting the original subspaces into an higher dimensional linear space does not increase the number of original non-overlapping dimensions. One then understands the role of measurement as a way to probe the differences between the classes, providing the degrees of freedom in order to distinguish between the signals from classes $C = 1$ and $C = 2$.

On the other hand, the measurement gain is a function of the exact geometry of the classes in the Gaussian mixture model. It increases with the ratio of the product of the nonzero eigenvalues of $\mathbf{\Phi}\left(\mathbf{\Sigma}_1 + \mathbf{\Sigma}_2\right)\mathbf{\Phi}^T$ to the product of the nonzero singular values of $\mathbf{\Phi}\mathbf{\Sigma}_1\mathbf{\Phi}^T$ and $\mathbf{\Phi}\mathbf{\Sigma}_2\mathbf{\Phi}^T$.

B. Nonzero-Mean Classes

The following Theorem now generalizes the description of the asymptotic behavior of the probability of misclassification from the zero-mean to the nonzero-mean, two-class compressive classification problem.

Theorem 3: Consider the measurement model in (1) where $\mathbf{x} \sim \mathcal{N}(\boldsymbol{\mu}_1, \mathbf{\Sigma}_1)$ with probability P_1 and $\mathbf{x} \sim \mathcal{N}(\boldsymbol{\mu}_2, \mathbf{\Sigma}_2)$ with probability $P_2 = 1 - P_1$. If

$$\mathbf{\Phi}(\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2) \notin \mathrm{im}(\mathbf{\Phi}(\mathbf{\Sigma}_1 + \mathbf{\Sigma}_2)\mathbf{\Phi}^T), \tag{21}$$

then the upper bound to the probability of misclassification in (14) decays exponentially with $1/\sigma^2$ as $\sigma^2 \to 0$; otherwise,

$$P_{err}^{UB} = \left(\frac{a \cdot g_m}{\sigma^2}\right)^{-d} + o\left(\left(\frac{1}{\sigma^2}\right)^{-d}\right), \quad \sigma^2 \to 0 \tag{22}$$

where a is a finite constant which depends only on the first term in (15), with $a = 1$ for $\boldsymbol{\mu}_1 = \boldsymbol{\mu}_2$ and $a > 1$ for $\boldsymbol{\mu}_1 \neq \boldsymbol{\mu}_2$, whereas g_m and d are as in Theorems 1 and 2.

Proof: The proof is presented in Appendix C. ■

The characterization embodied in Theorem 3 illustrates that the asymptotic behavior of the upper bound of the error probability for classes with nonzero-mean can be radically different from that for classes with zero-mean. The differences in behavior trace back to the fact that $M > r_{\mathbf{\Sigma}_{12}}$ represents a necessary condition for condition (21) to hold (see Appendix C). In the nonzero-mean case, choosing $M > r_{\mathbf{\Sigma}_{12}}$ (provided that condition (21) also holds) leads to a diversity-order $d = \infty$; in contrast, in the zero-mean case choosing $M > r_{\mathbf{\Sigma}_{12}}$ does not affect the diversity-order and only affects the measurement gain. In contrast, letting $M \leq r_{\mathbf{\Sigma}_{12}}$ induces the same diversity-order both for nonzero-mean and zero-mean classes; the presence of the nonzero-mean here may only impact the measurement gain since $a \geq 1$.

In other terms, we can argue that we will not achieve infinite diversity-order if the projection (according to $\mathbf{\Phi}$) of the difference $\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2$ lies in the direct sum of the projected sub-spaces corresponding to the two classes. In

this case, as expected, the fact that the two classes have a nonzero-mean does not change the diversity-order as the mean vectors do not provide a way to separate the projected sub-spaces corresponding to the two classes. On the other hand, when the projection of the difference $\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2$ does not lie in the direct sum of the projected sub-spaces corresponding to the two classes, the upper bound on the misclassification probability decays exponentially with $1/\sigma^2$ when $\sigma^2 \to 0$. Geometrically, this result associated with infinite diversity reflects the fact that, when embedded in a higher dimensional space ($\mathbb{R}^M$ in our cases), the affine spaces spanned by the linear transformation of the signals in the two classes do not intersect.

IV. Multiple-Class Compressive Classification with Random Measurements in the Low-Noise Regime

We now consider the characterization of the performance of a multiple-class compressive classification problem, where $L \geq 3$, with random measurements. The generalization of the two-class results to the multiple-class case is possible by using the union bound in conjunction with the two-class Bhattacharyya bound that leads to the upper bound to the probability of misclassification given by [42]:

$$P_{err}^{UB} = \sum_{i=1}^{L} \sum_{\substack{j=1 \\ j \neq i}}^{L} P_i \; e^{-K_{ij}}, \tag{23}$$

where K_{ij} is also given by (15).

The fact that the form of the upper bound in (23) is akin to the form of the upper bound in (14), involving in addition only the various pairwise misclassification terms that capture the interaction between the different classes, leads to the immediate generalization of the results encapsulated in the previous Theorems. In fact, we can easily conclude from (3) and (23) that the diversity-order for the multiple-class classification problem is given by

$$d = \min_{\substack{i,j \\ j \neq i}} d(i,j), \tag{24}$$

where $d\,(i,j)$ represents the diversity-order of a given pair of classes $C = i$ and $C = j$ in the two-class classification problem[6]. In addition, we can also conclude from (4) and (23) that the measurement gain for the multiple-class classification problem is given by

$$g_m = \left[\sum_{(i,j) \in \mathcal{S}_d} P_i 2^{r_{ij}} \left[\frac{v_{ij}}{\sqrt{v_i v_j}} \right]^{-\frac{1}{2}} \right]^{-\frac{1}{d}} \tag{25}$$

where $\mathcal{S}_d$ is the set of pairs of indexes corresponding to pairs of classes with minimum diversity-order, that is, $\mathcal{S}_d = \{(i,j) : i \neq j, d(i,j) = d\}$.

In particular, we can argue that the upper bound to the misclassification probability will exhibit an error floor if at least one of the pairwise misclassification probabilities also exhibits an error floor. Conversely, the upper bound to the misclassification probability (and the true misclassification probability) will tend to zero as σ^2 tends to zero

[6]Note that the diversity-order $d(i,j)$ does not depend on the exact value of the a priori probability of the classes P_i, P_j, provided that $P_i, P_j > 0$.

if all the pairwise upper bounds to the misclassification probabilities also tend to zero. We can also understand how the performance is affected by the geometry of the problem via (24), (25) together with Theorems 1, 2 and 3.

V. Measurements Design in the Low-Noise Regime

It is also of interest to investigate how one can further improve performance by using designed measurements in *lieu* of the conventional random ones in the low-noise regime. In particular, we investigate this question by posing a measurement design problem where the objective is to maximize the diversity-order subject to a given measurement budget, i.e.

$$\max_{\mathbf{\Phi}} d\left(\mathbf{\Phi}\right), \tag{26}$$

subject to:

$$\mathrm{rank}\left(\mathbf{\Phi}\right) \leq M. \tag{27}$$

where we express explicitly the diversity-order in terms of the measurement matrix. [7]

In the sequel, and in line with the previous analytical procedure, we will consider measurement designs for a two-class scenario followed by measurement designs for the multiple-class case. For the two-class problem, it is possible to solve the optimization problem in (26)–(27); for multiple-class problems, it does not seem possible to conceive a closed-form solution to such optimization problem – therefore, we only put forth an, in general suboptimal, algorithm which is inspired by the optimal solution of the two-class scenario: this algorithm attempts at maximizing the diversity-order while satisfying the measurement budget constraint.

A. Two-Class Case

We first consider kernel designs for two-class compressive classification problems.

1) Zero-Mean Classes: The following Theorem defines the kernel design that solves the optimization problem in (26) and (27) for the compressive classification of two zero-mean classes.

Theorem 4: Consider the measurement model in (1) where $\mathbf{x} \sim \mathcal{N}(\mathbf{0}, \mathbf{\Sigma}_1)$ with probability P_1 and $\mathbf{x} \sim \mathcal{N}(\mathbf{0}, \mathbf{\Sigma}_2)$ with probability $P_2 = 1 - P_1$. Assume that the measurement budget is such that $M \geq NO_{Dim}$. Then, the maximum possible diversity-order is given by:

$$d_{max} = \frac{1}{4} NO_{Dim} \tag{28}$$

which is achieved by a measurement matrix design that obeys the following necessary and sufficient condition:

$$2r_{12} - r_1 - r_2 = NO_{Dim}. \tag{29}$$

[7] In view of the fact that we use the diversity-order and the measurement gain as a proxy to measure performance in the low-noise regime, we could also conceive measurement design problems where one would first define the set of kernels that maximize the diversity-order subject to the measurement budget and one would then define – out of the diversity-order maximizing designs - the kernel that maximizes the measurement gain. One could also impose other additional constraints, such as an average power constraint. Our designs also respect this constraint.

A measurement matrix $\mathbf{\Phi}$ that achieves the maximum diversity-order is

$$\mathbf{\Phi} = \left[\mathbf{v}_1, \mathbf{v}_2, \ldots, \mathbf{v}_{n_{\mathbf{\Sigma}_1}}, \mathbf{w}_1, \mathbf{w}_2, \ldots, \mathbf{w}_{n_{\mathbf{\Sigma}_2}}\right]^T, \tag{30}$$

where the sets of vectors $\left[\boldsymbol{u}_1, \ldots, \boldsymbol{u}_{n_{12}}\right]$, $\left[\boldsymbol{u}_1, \ldots, \boldsymbol{u}_{n_{12}}, \boldsymbol{v}_1, \ldots, \boldsymbol{v}_{n_{\mathbf{\Sigma}_1}}\right]$, $\left[\boldsymbol{u}_1, \ldots, \boldsymbol{u}_{n_{12}}, \boldsymbol{w}_1, \ldots, \boldsymbol{w}_{n_{\mathbf{\Sigma}_2}}\right]$, $\boldsymbol{u}_i, \boldsymbol{v}_i, \boldsymbol{w}_i \in \mathbb{R}^N$, constitute an orthonormal basis of the linear spaces $\mathrm{Null}\left(\mathbf{\Sigma}_1\right) \bigcap \mathrm{Null}\left(\mathbf{\Sigma}_2\right)$, $\mathrm{Null}\left(\mathbf{\Sigma}_1\right)$ and $\mathrm{Null}\left(\mathbf{\Sigma}_2\right)$, respectively, and $n_{12} = N - r_{\mathbf{\Sigma}_{12}}$, $n_{\mathbf{\Sigma}_1} = N - n_{12} - r_{\mathbf{\Sigma}_1}$ and $n_{\mathbf{\Sigma}_2} = N - n_{12} - r_{\mathbf{\Sigma}_2}$.

Assume now that the measurement budget is such that $M < NO_{Dim}$. Then, the maximum possible diversity-order is given by:

$$d = \frac{1}{4}M \tag{31}$$

which is achieved, if and only if, the measurement matrix design is such that $r_{12} = M$ and $r_1 + r_2 = r_{12}$. A measurement matrix $\mathbf{\Phi}$ that achieves such a diversity-order can be obtained from the measurement matrix $\mathbf{\Phi}$ in (30), by choosing arbitrarily only M out of the $n_{\mathbf{\Sigma}_1} + n_{\mathbf{\Sigma}_2}$ row vectors.

Proof: The proof is presented in Appendix D. ■

This Theorem encapsulates key operational features associated with optimized measurements:

- For a sufficient measurement budget, the Theorem reveals that the maximum diversity-order is directly related to the number of non-overlapping dimensions between the two classes. The measurements design that achieves such maximum diversity-order in (28) is then required to satisfy $2r_{12} - r_1 - r_2 = NO_{Dim}$ that implies – as shown in Appendix D – that i) it measures all the non-overlapping dimensions, i.e. $M \geq NO_{Dim}$ and ii) it measures all the dimensions in each class that are not contained in the intersection of the corresponding sub-spaces, i.e. $r_1 \geq r_{\mathbf{\Sigma}_{12}} - r_{\mathbf{\Sigma}_2}$ and $r_2 \geq r_{\mathbf{\Sigma}_{12}} - r_{\mathbf{\Sigma}_1}$. In contrast, for an insufficient measurement budget, the Theorem reveals that maximum diversity-order is related instead to the number of available measurements. The measurement designs that achieve this maximal diversity-order in (31) now only probe a limited number of the non-overlapping dimensions. Therefore we can argue that optimal measurement employs features associated with linear subspaces that are contained in the spaces spanned by the individual classes but not contained in their intersection.
- In particular, it follows from Theorems 2 and 4 that a designed kernel can offer marked improvements over a random one in the low-noise regime. These include:
  i) the ability to achieve perfect separation of the projected signals with a single measurement – with a random measurement kernel, according to Theorem 2, we require $M > r_{\mathbf{\Sigma}_1}$ in order to drive the (upper bound to the) probability of misclassification to zero as the noise level approaches zero;
  ii) the ability to achieve the maximum diversity-order with a lower number of measurements – with a random measurement kernel, according to Theorem 2, we require $M = r_{\mathbf{\Sigma}_{12}}$ in order to extract the maximum diversity-order, but with a designed measurement kernel we only require $M = NO_{Dim} \leq r_{\mathbf{\Sigma}_{12}}$;

2) Nonzero-Mean Classes: The following Theorem now defines the kernel design that solves the optimization problem in (26) and (27) for the compressive classification of two nonzero-mean classes.

Theorem 5: Consider the measurement model in (1) where $\mathbf{x} \sim \mathcal{N}(\boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1)$ with probability P_1 and $\mathbf{x} \sim \mathcal{N}(\boldsymbol{\mu}_2, \boldsymbol{\Sigma}_2)$ with probability $P_2 = 1 - P_1$. Assume that:

$$(\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2) \notin \operatorname{im}(\boldsymbol{\Sigma}_1 + \boldsymbol{\Sigma}_2). \tag{32}$$

Then, the maximum diversity-order is $d = \infty$ and a matrix design that achieves such a diversity-order is:

$$\boldsymbol{\Phi} = \left[\boldsymbol{\phi}^T\right], \tag{33}$$

where $\boldsymbol{\phi}$ can be any vector that belongs to $\operatorname{Null}(\boldsymbol{\Sigma}_1 + \boldsymbol{\Sigma}_2)$.

Assume now that:

$$(\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2) \in \operatorname{im}(\boldsymbol{\Sigma}_1 + \boldsymbol{\Sigma}_2). \tag{34}$$

Then, the maximum diversity-order is given by (28) or (31), depending on the number of available measurements and a matrix design that achieves such a diversity-order is given by Theorem 4.

Proof: The proof is presented in Appendix E. ■

The most important feature associated with this Theorem relates to the fact that, under certain conditions associated with the geometry of the classification problem, it is possible to attain a diversity-order $d = \infty$, or exponential decay of the error probability, by taking a single measurement. This feature, which is also unique to nonzero-mean classes, also bears witness to the advantage of kernel design: recall that random measurement kernel requires $M > r_{\boldsymbol{\Sigma}_{12}}$ in order to achieve a diversity-order $d = \infty$ (see Theorem 3). Note also that the existence of a $\boldsymbol{\phi} \in \operatorname{Null}(\boldsymbol{\Sigma}_1 + \boldsymbol{\Sigma}_2)$, which is used to construct the measurement kernel in (33), is guaranteed by the condition in (32).

B. Multiple-Class Case

We now consider kernel designs for multiple-class compressive classification problems. In particular, we propose, possibly suboptimal, algorithmic approaches, inspired by the two-class case designs in Theorems 4 and 5, that attempt to achieve the maximum diversity-order while satisfying the constraint on the number of measurements. We leverage the fact that the diversity-order associated with a multiple-class classification problem corresponds to the lowest of the diversity-orders of the pairwise classification problems (see (24)).

1) Zero-Mean Classes: Table I puts forth an algorithmic approach to design projections for multiple-class classification problems with zero-mean classes. The algorithm consists of four main steps. In the first step, we identify the pair of classes associated with the minimum value of the diversity-order (for the two-class compressive classification problem). In the second step we construct a set of measurement matrices $\boldsymbol{\Phi}_{ij}, \forall i, j$ as in Theorem 4 for all possible pairs of classes, consisting of a number of measurements necessary to achieve the minimum value of the diversity-order determined in step 1. In the third step we build a measurement matrix by concatenating

Table I:

Algorithmic approach to maximize the diversity-order subject to a measurement budget in a multiple-class classification problem with zero-mean classes.

Input : Number of classes L, input covariance matrices $\mathbf{\Sigma}_i$, $i = 1, \cdots, L$, number of measurements M.

Output: Measurement matrix $\mathbf{\Phi}$.

Step 1 – Determine $(i^*, j^*) = \arg\min_{\substack{i,j \\ i \neq j}} \; NO_{Dim}(i,j)$, where $NO_{Dim}(i,j)$ is as expressed in (11);

Step 2 – **for** $i = 1, \ldots, L-1$

  – **for** $j = i+1, \ldots, L$

    – Construct $\mathbf{\Phi}_{ij} \in \mathbb{R}^{NO_{Dim}(i^*,j^*) \times N}$ according to Theorem 4 applied to classes i and j;

Step 3 – $\mathbf{\Phi} \leftarrow \left[\mathbf{\Phi}_{12}^T, \mathbf{\Phi}_{13}^T, \ldots, \mathbf{\Phi}_{(L-1)L}^T\right]^T$;

Step 4 – Compute the rank $r_{\mathbf{\Phi}} = \mathrm{rank}(\mathbf{\Phi})$:

– **if** $r_{\mathbf{\Phi}} \leq M$

  – Return the matrix obtained by selecting the nonzero rows of the row echelon form of $\mathbf{\Phi}$;

– **else**

  – $\mathbf{\Phi}_{ij} \leftarrow \left[\mathbf{I}_{M_{ij}-1}, \mathbf{0}_{(M_{ij}-1)\times 1}\right] \mathbf{\Phi}_{ij}$, where M_{ij} denotes the number of rows of $\mathbf{\Phi}_{ij}$;

  – Go to Step 3;

all the measurements designed for the different pairs of classes. Then, in the fourth step we compute the rank of such matrix. If the rank is less than or equal to the number of available measurements, then the algorithm returns the matrix obtained by picking the nonzero rows of the row echelon form of matrix $\mathbf{\Phi}$. Otherwise, for all pair of classes, we erase the last measurement from the corresponding matrix $\mathbf{\Phi}_{ij}$ and iterate again from step 3 until the rank of $\mathbf{\Phi}$ satisfies the measurement constraint.

Note that the construction in Algorithm 1 satisfies the measurement budget constraint, but it does not necessarily lead to the global optimum solution of the problem in (26)–(27), as different choices of measurements in step 2 and step 4 might lead to different values of the diversity-order.[8] However, observe that the matrix constructed in step 2 achieves the maximum diversity-order, and each iteration of step 4 decreases the achieved diversity-order of a factor at most equal to $1/4$. Finally, we underline the fact that Algorithm 1 might eventually output an empty measurement matrix. This can happen if, after $NO_{Dim}(i^*, j^*) - 1$ iterations of step 4, the single measurements corresponding to the $L(L-1)/2$ pairs of classes still span a linear space of dimension greater than M. In order to avoid that the algorithm outputs such an empty matrix, we could propose to delete measurements only from a single randomly selected matrix $\mathbf{\Phi}_{ij}$ or some randomly selected subset of matrices $\mathbf{\Phi}_{ij}$ *in lieu* of a single measurement from all such matrices: this procedure would never result in the output of a null matrix, but could result in an error

[8]Note that, in step 2, we are arbitrarily choosing $NO_{Dim}(i^*, j^*)$ linearly independent measurements from a linear space of dimension $NO_{Dim}(i,j)$. Moreover, in step 4 we could decide to delete any row of $\mathbf{\Phi}_{ij}$ instead of the last one.

Table II:

Algorithmic approach to maximize the diversity-order subject to a measurement budget in the multiple-class classification problem with nonzero-mean classes.

Input : Number of classes L, input covariance matrices $\mathbf{\Sigma}_i$ and mean vectors $\boldsymbol{\mu}_i$, $i = 1, \cdots, L$, number of measurements M.

Output: Measurement matrix $\mathbf{\Phi}$.

Step 1 – **for** $i = 1, \ldots, L-1$

- **for** $j = i+1, \ldots, L$
  - **if** $(\boldsymbol{\mu}_i - \boldsymbol{\mu}_j) \notin \mathrm{im}\,(\mathbf{\Sigma}_i + \mathbf{\Sigma}_j)$
    - Choose $\boldsymbol{\phi}_{ij} \in \mathrm{Null}\,(\mathbf{\Sigma}_i + \mathbf{\Sigma}_j)$ as in Theorem 5;
  - **else**
    - Go to Step 4;

Step 2 – $\mathbf{\Phi} \leftarrow \left[\boldsymbol{\phi}_{12}, \boldsymbol{\phi}_{13}, \ldots, \phi_{(L-1)L}\right]^T$;

Step 3 – Compute the rank $r_{\mathbf{\Phi}} = \mathrm{rank}(\mathbf{\Phi})$:

- **if** $r_{\mathbf{\Phi}} \leq M$
  - Return the matrix obtained by selecting the nonzero rows of the row echelon form of $\mathbf{\Phi}$;
- **else**
  - Go to Step 4;

Step 4 – Construct $\mathbf{\Phi}$ according to Algorithm 1;

floor in the upper bound of the misclassification probability. Despite this possible issues, numerical results show that matrices designed according to this algorithm appear to lead to very good performance.

2) Nonzero-Mean Classes: The algorithmic approach to design projections for multiple-class classification problems with nonzero-mean classes is shown in Table II.

The algorithm also consists of four steps. In step 1, we check if $(\boldsymbol{\mu}_i - \boldsymbol{\mu}_j) \notin \mathrm{im}\,(\mathbf{\Sigma}_i + \mathbf{\Sigma}_j)$ holds for every pair of classes. In this case, we can construct the measurement kernel by picking vectors from the null spaces $\mathrm{Null}\,(\mathbf{\Sigma}_i + \mathbf{\Sigma}_j)$ for every pair of classes, thus constituting a measurement matrix $\mathbf{\Phi}$. In step 3 we check if the rank of such matrix satisfies the measurement constraint. If this is the case, the algorithm returns the matrix obtained by picking the nonzero rows of the row echelon form of matrix $\mathbf{\Phi}$. Otherwise, or if there is a pair of classes for which $(\boldsymbol{\mu}_i - \boldsymbol{\mu}_j) \in \mathrm{im}\,(\mathbf{\Sigma}_i + \mathbf{\Sigma}_j)$, then, it is not possible to achieve infinite diversity-order with M measurements extracted from $\mathbf{\Phi}$ and we construct the measurement kernel according to Algorithm 1.

VI. High-Noise Regime

It is also of interest to briefly contrast the behavior of the upper bound to the misclassification probability in the low-noise regime ($\sigma^2 \to 0$) to the high-noise regime ($\sigma^2 \to \infty$). For example, for a two-class compressive classification problem where $\mathbf{x} \sim \mathcal{N}(\boldsymbol{\mu}_1, \mathbf{\Sigma}_1)$ with probability P_1 and $\mathbf{x} \sim \mathcal{N}(\boldsymbol{\mu}_2, \mathbf{\Sigma}_2)$ with probability $P_2 = 1 - P_1$,

this can be done by expanding the Bhattacharyya upper bound to the misclassification probability, in (14) and (15), as follows[9]:

$$P_{err}^{UB}\left(\frac{1}{\sigma^2}\right)=P_{err}^{UB}\left(0\right)+\frac{d\ P_{err}^{UB}\left(\frac{1}{\sigma^2}\right)}{d\ \frac{1}{\sigma^2}}\Big|_{\frac{1}{\sigma^2}=0}\frac{1}{\sigma^2}+\frac{d^2\ P_{err}^{UB}\left(\frac{1}{\sigma^2}\right)}{d\ \left(\frac{1}{\sigma^2}\right)^2}\Big|_{\frac{1}{\sigma^2}=0}\frac{\left(\frac{1}{\sigma^2}\right)^2}{2}+o\left(\left(\frac{1}{\sigma^2}\right)^2\right). \tag{35}$$

Then, for zero-mean classes, the expansion of the upper bound to the probability of misclassification reduces to:

$$P_{err}^{UB}\left(\frac{1}{\sigma^2}\right)=\sqrt{P_1P_2}\ \left[1+\frac{1}{4}\ A\ \left(\frac{1}{\sigma^2}\right)^2\right]+o\left(\left(\frac{1}{\sigma^2}\right)^2\right), \tag{36}$$

where:

$$\begin{aligned}A=&\mathrm{tr}\left[\left(\frac{\mathbf{\Phi}\left(\mathbf{\Sigma}_1+\mathbf{\Sigma}_2\right)\mathbf{\Phi}^T}{2}\right)^2\right]-\frac{1}{2}\mathrm{tr}\left[\left(\mathbf{\Phi}\mathbf{\Sigma}_1\mathbf{\Phi}^T\right)^2\right]-\frac{1}{2}\mathrm{tr}\left[\left(\mathbf{\Phi}\mathbf{\Sigma}_2\mathbf{\Phi}^T\right)^2\right]\\&+\mathrm{tr}\left[\mathbf{\Phi}\mathbf{\Sigma}_1\mathbf{\Phi}^T\right]\mathrm{tr}\left[\mathbf{\Phi}\mathbf{\Sigma}_2\mathbf{\Phi}^T\right]-\mathrm{tr}^2\left[\frac{\mathbf{\Phi}\left(\mathbf{\Sigma}_1+\mathbf{\Sigma}_2\right)\mathbf{\Phi}^T}{2}\right],\end{aligned} \tag{37}$$

and for nonzero-mean classes, it reduces to:

$$P_{err}^{UB}\left(\frac{1}{\sigma^2}\right)=\sqrt{P_1P_2}\ \left[1-\frac{1}{8}\left\|\mathbf{\Phi}\left(\boldsymbol{\mu}_1-\boldsymbol{\mu}_2\right)\right\|^2\frac{1}{\sigma^2}\right]+o\left(\frac{1}{\sigma^2}\right). \tag{38}$$

Note that the behavior of the upper bound to the misclassification probability for zero-mean classes is fundamentally different from that for nonzero-mean classes. In particular, the first-order term in (35) is always equal to zero for zero-mean classes and is nonzero for nonzero-mean classes, implying that the upper bound to the misclassification probability decays faster in the later case as $1/\sigma^2\to 0$. In addition, the behavior of the high-noise expansions is also fundamentally different from that of the low-noise ones. In the low-noise case, we can approximate the upper bound to the misclassification probability by a line defined via its slope (the diversity-order) and its offset (the measurement gain) on a log-log scale. Of course, its slope can occasionally be infinite as per Theorem 3. In contrast, in the high-noise case we approximate the upper bound to the misclassification probability with a line (for nonzero-mean classes) or with a parabola (for zero-mean classes) on a linear, rather than logarithmic, scale.

We can also derive further insight by studying the behavior of the average, with respect to the measurement matrix, of the upper bounds to the misclassification probability by assuming that the measurement matrix is such that its elements are drawn i.i.d. from a zero-mean unit-variance Gaussian distribution. In particular, it is simple to establish that the average of the upper bound to the misclassification probability behaves as follows:

$$\bar{P}_{err}^{UB}\left(\frac{1}{\sigma^2}\right)=\mathcal{E}\left\{P_{err}^{UB}\left(\frac{1}{\sigma^2}\right)\right\}=\sqrt{P_1P_2}\ \left[1+\frac{1}{4}\ \mathcal{E}\left\{A\right\}\ \left(\frac{1}{\sigma^2}\right)^2\right]+o\left(\left(\frac{1}{\sigma^2}\right)^2\right), \tag{39}$$

for zero-mean classes and:

$$\bar{P}_{err}^{UB}\left(\frac{1}{\sigma^2}\right)=\mathcal{E}\left\{P_{err}^{UB}\left(\frac{1}{\sigma^2}\right)\right\}=\sqrt{P_1P_2}\ \left[1-\frac{1}{8}\ \mathcal{E}\left\{\left\|\mathbf{\Phi}\left(\boldsymbol{\mu}_1-\boldsymbol{\mu}_2\right)\right\|^2\right\}\frac{1}{\sigma^2}\right]+o\left(\frac{1}{\sigma^2}\right), \tag{40}$$

for nonzero-mean classes. Via random matrix theory, we can further calculate:

$$\mathcal{E}\left\{\mathrm{tr}\left[\mathbf{\Phi}\mathbf{\Sigma}_1\mathbf{\Phi}^T\right]\mathrm{tr}\left[\mathbf{\Phi}\mathbf{\Sigma}_2\mathbf{\Phi}^T\right]\right\}=M\left(\mathrm{tr}\left[\mathbf{\Sigma}_1\right]\mathrm{tr}\left[\mathbf{\Sigma}_2\right]+2\mathrm{tr}\left[\mathbf{\Sigma}_1\mathbf{\Sigma}_2\right]\right)+M\left(M-1\right)\mathrm{tr}\left[\mathbf{\Sigma}_1\right]\mathrm{tr}\left[\mathbf{\Sigma}_2\right] \tag{41}$$

[9]The expansion and ensuing discussion also generalize immediately to multiple-class problems via the usual union bound arguments.

$$\mathcal{E}\left\{\mathrm{tr}^2\left[\frac{\mathbf{\Phi}\left(\mathbf{\Sigma}_1+\mathbf{\Sigma}_2\right)\mathbf{\Phi}^T}{2}\right]\right\}=M\left\{\mathrm{tr}^2\left[\frac{\left(\mathbf{\Sigma}_1+\mathbf{\Sigma}_2\right)}{2}\right]+2\mathrm{tr}\left[\left(\frac{\left(\mathbf{\Sigma}_1+\mathbf{\Sigma}_2\right)}{2}\right)^2\right]\right\}$$
$$+M\left(M-1\right)\mathrm{tr}^2\left[\frac{\left(\mathbf{\Sigma}_1+\mathbf{\Sigma}_2\right)}{2}\right] \tag{42}$$

$$\mathcal{E}\left\{\mathrm{tr}\left[\left(\mathbf{\Phi}\mathbf{\Sigma}_1\mathbf{\Phi}^T\right)^2\right]\right\}=M\left[\sum_{j=1}^{N}m_4{\lambda_{\mathbf{\Sigma}_1}}_j^2+\sum_{\substack{j,k=1\\j\neq k}}^{N}m_2\lambda_{\mathbf{\Sigma}_1 j}\lambda_{\mathbf{\Sigma}_1 k}\right]+M\left(M-1\right)\sum_{j=1}^{N}m_2{\lambda_{\mathbf{\Sigma}_1}}_j^2 \tag{43}$$

$$\mathcal{E}\left\{\mathrm{tr}\left[\left(\mathbf{\Phi}\mathbf{\Sigma}_2\mathbf{\Phi}^T\right)^2\right]\right\}=M\left[\sum_{j=1}^{N}m_4{\lambda_{\mathbf{\Sigma}_2}}_j^2+\sum_{\substack{j,k=1\\j\neq k}}^{N}m_2\lambda_{\mathbf{\Sigma}_2 j}\lambda_{\mathbf{\Sigma}_2 k}\right]+M\left(M-1\right)\sum_{j=1}^{N}m_2{\lambda_{\mathbf{\Sigma}_2}}_j^2 \tag{44}$$

$$\mathcal{E}\left\{\mathrm{tr}\left[\left(\frac{\mathbf{\Phi}\left(\mathbf{\Sigma}_1+\mathbf{\Sigma}_2\right)\mathbf{\Phi}^T}{2}\right)^2\right]\right\}=M\left[\sum_{j=1}^{N}m_4{\lambda_{\mathbf{\Sigma}_{12}}}_j^2+\sum_{\substack{j,k=1\\j\neq k}}^{N}m_2\lambda_{\mathbf{\Sigma}_{12}j}\lambda_{\mathbf{\Sigma}_{12}k}\right]+M\left(M-1\right)\sum_{j=1}^{N}m_2{\lambda_{\mathbf{\Sigma}_{12}}}_j^2 \tag{45}$$

$$\mathcal{E}\left\{\|\mathbf{\Phi}\left(\boldsymbol{\mu}_1-\boldsymbol{\mu}_2\right)\|^2\right\}=M\|\boldsymbol{\mu}_1-\boldsymbol{\mu}_2\|^2 \tag{46}$$

where $\lambda_{\mathbf{\Sigma}_1 j}$, $j=1,\ldots,N$ are the eigenvalues of the matrix $\mathbf{\Sigma}_1$, $\lambda_{\mathbf{\Sigma}_2 j}$, $j=1,\ldots,N$ are the eigenvalues of the matrix $\mathbf{\Sigma}_2$, $\lambda_{\mathbf{\Sigma}_{12} j}$, $j=1,\ldots,N$ are the eigenvalues of the matrix $\frac{(\mathbf{\Sigma}_1+\mathbf{\Sigma}_2)}{2}$, and m_2 and m_4 are the second and fourth-order moments, respectively, of a zero-mean and unit variance Gaussian random variable.

Therefore, one observes that in the nonzero-mean case the behavior of the average value of the upper bound to the misclassification probability depends only on the number of measurements and the means of the classes – moreover, the higher the number of measurements the lower the average upper bound; in contrast, in the zero-mean case the average value of the upper bound to the probability of the misclassification depends in a more intricate manner on the number of measurements and the source covariances via the eigenvalues of the matrices $\mathbf{\Sigma}_1$, $\mathbf{\Sigma}_2$ and $\frac{(\mathbf{\Sigma}_1+\mathbf{\Sigma}_2)}{2}$.

It is more difficult to establish how the upper bound to the misclassification probability, in the high-noise regime, behaves in the presence of optimized measurements (though equations (36) and (38) offer a means to carry out numerical optimizations). However, we can easily see that for a compressive classification problem with nonzero-mean classes the design which minimizes the first-order expansion of the upper bound to the probability of misclassification in the high-noise regime is obtained by aligning the measurements with the vector $(\boldsymbol{\mu}_1-\boldsymbol{\mu}_2)$, i.e.:

$$\mathbf{\Phi}=\alpha\left(\boldsymbol{\mu}_1-\boldsymbol{\mu}_2\right)^T, \tag{47}$$

where the scalar $\alpha\neq 0$ determines the norm of the measurement vector $\mathbf{\Phi}$.

VII. Numerical Results

We now present a series of results that illustrate the main operational features associated with the compressive classification of a mixture of Gaussians. In particular, we also compare the behavior of the upper bound to the probability of misclassification to the behavior of the true probability of misclassification, in order to determine whether the previous theoretical results are aligned with real ones.

A. Random Measurements

We first consider a compressive classification problem with two zero-mean classes. The two classes are such that $\boldsymbol{\mu}_1 = \boldsymbol{\mu}_2 = 0$, $\boldsymbol{\Sigma}_1 = \mathbf{U}\text{diag}\,(1,1,0,0,0,0)\,\mathbf{U}^T$ and $\boldsymbol{\Sigma}_2 = \mathbf{U}\text{diag}\,(0,1,1,1,0,0)\,\mathbf{U}^T$, where $\mathbf{U}$ is a randomly generated unitary matrix. Therefore, $r_{\boldsymbol{\Sigma}_1} = \text{rank}\,(\boldsymbol{\Sigma}_1) = 2$, $r_{\boldsymbol{\Sigma}_2} = \text{rank}\,(\boldsymbol{\Sigma}_2) = 3$, $r_{\boldsymbol{\Sigma}_{12}} = \text{rank}\,(\boldsymbol{\Sigma}_1 + \boldsymbol{\Sigma}_2) = 4$ and $N = 6$. The sensing matrix is such that:

$$\boldsymbol{\Phi} = \frac{M}{\text{tr}\left(\boldsymbol{\Phi}'\boldsymbol{\Phi}'^T\right)}\boldsymbol{\Phi}', \tag{48}$$

where $\boldsymbol{\Phi}'$ has i.i.d. zero-mean unit-variance Gaussian entries. Note that this scenario is such that the sub-spaces corresponding to the two classes do not overlap completely and the number of non-overlapping dimensions in (11) is equal to 3. Therefore, accordingly to Theorems 1 and 2 it is possible to drive the upper bound to the misclassification probability to zero at low-noise levels.

Figure 1(a) shows that for $M = 2$ the upper bound exhibits an error floor and for $M > 2$ the upper bound tends to zero as the noise level also tends to zero, in accordance with Theorem 2. Figure 1(a) also shows that the increase of M from 2 to 3 or 4 results in the increase in the diversity-order but $M > r_{\boldsymbol{\Sigma}_{12}}$ does not result in further increases in diversity but only in measurement gain – this is also consistent with Theorem 2. We can also observe that the upper bound to the probability of misclassification is able to capture the behavior of the true probability of misclassification: it captures the presence or absence of an error floor and (except for $M = 2$) it also captures closely the diversity-order and increases in measurement gain in the true error probability.

We now consider a compressive classification problem with two nonzero-mean classes. Here, the two classes are such that $\boldsymbol{\mu}_1 \neq \boldsymbol{\mu}_2$ and $\boldsymbol{\Sigma}_1 = \boldsymbol{\Sigma}_2 = \mathbf{U}\text{diag}\,(1,1,0,0,0,0)\,\mathbf{U}^T$, where $\mathbf{U}$ is also a randomly generated unitary matrix. Then, $r_{\boldsymbol{\Sigma}_1} = \text{rank}\,(\boldsymbol{\Sigma}_1) = 2$, $r_{\boldsymbol{\Sigma}_2} = \text{rank}\,(\boldsymbol{\Sigma}_2) = 2$, $r_{\boldsymbol{\Sigma}_{12}} = \text{rank}\,(\boldsymbol{\Sigma}_1 + \boldsymbol{\Sigma}_2) = 2$ and $N = 6$. The sensing matrix is also generated randomly with i.i.d. zero-mean unit-variance Gaussian entries and normalized as in the previous scenario. Note now that this scenario is such that the affine spaces corresponding to the two classes are parallel, differing only by a translation determined by the vector $\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2$.

Figure 1(b) – in line with Theorem 3 – shows that when $M \leq r_{\boldsymbol{\Sigma}_{12}} = 2$ the upper bound exhibits an error floor and when $M > r_{\boldsymbol{\Sigma}_{12}}$ the upper bounds to the misclassification probability tends to zero exponentially as the noise level also tends to zero. Once again, we can observe that the behavior of the upper bound to the probability of misclassification is consistent with the behavior of the true probability of misclassification.

We now turn the attention to a multi-class scenario where the number of classes is $L = 4$, $\boldsymbol{\mu}_1 = \boldsymbol{\mu}_2 = \boldsymbol{\mu}_3 = \boldsymbol{\mu}_4 = 0$, $\boldsymbol{\Sigma}_1 = \mathbf{U}\text{diag}\,(1,1,0,0,0,0)\,\mathbf{U}^T$, $\boldsymbol{\Sigma}_2 = \mathbf{U}\text{diag}\,(0,1,1,1,0,0)\,\mathbf{U}^T$, $\boldsymbol{\Sigma}_3 = \mathbf{U}\text{diag}\,(0,0,1,1,1,0)\,\mathbf{U}^T$ and

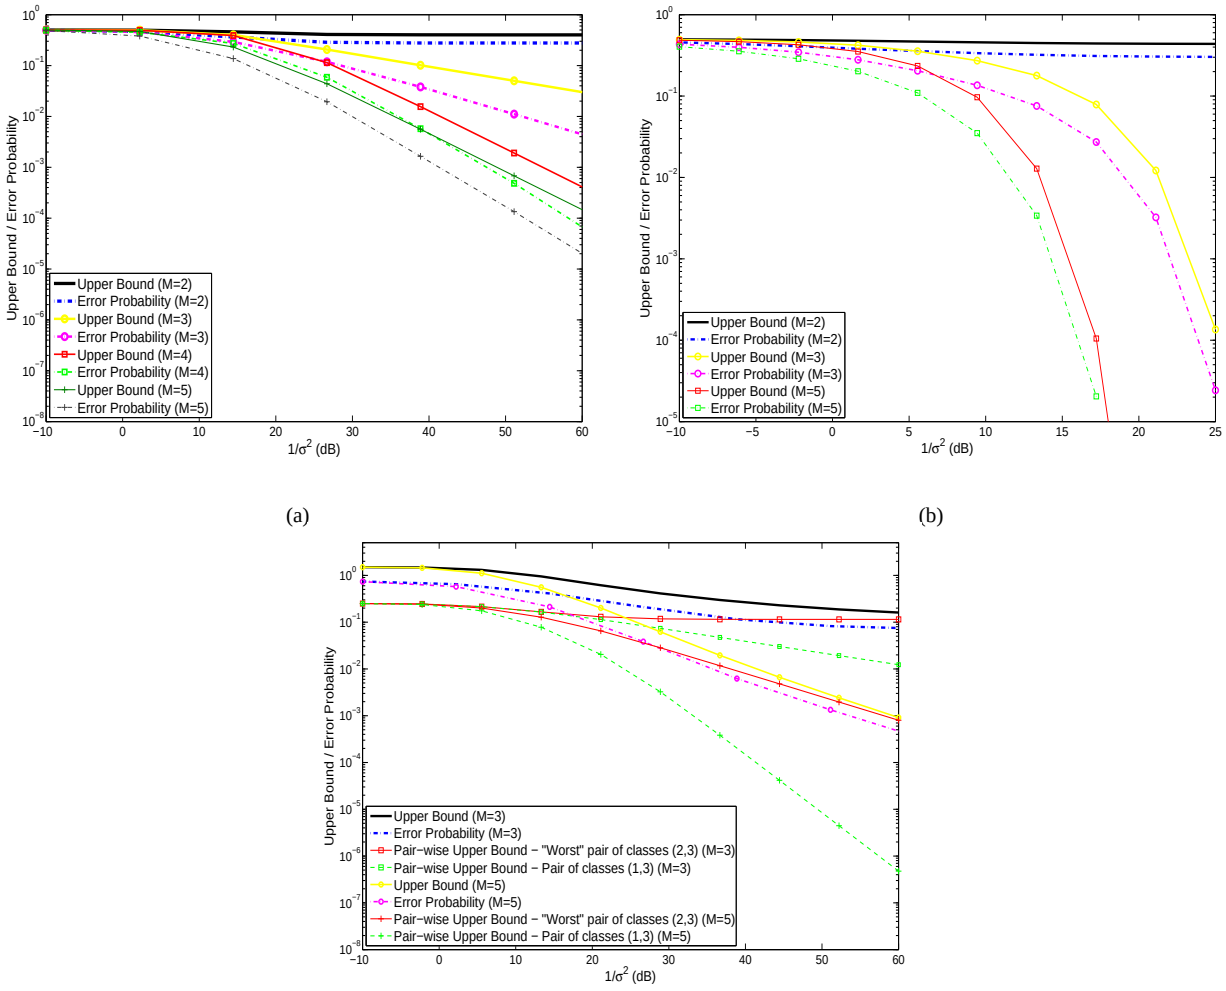
Upper Bound / Error Probability
1/σ² (dB)
Upper Bound (M=2)
Error Probability (M=2)
Upper Bound (M=3)
Error Probability (M=3)
Upper Bound (M=4)
Error Probability (M=4)
Upper Bound (M=5)
Error Probability (M=5)
(a)
(b)
Pair−wise Upper Bound − "Worst" pair of classes (2,3) (M=3)
Pair−wise Upper Bound − Pair of classes (1,3) (M=3)
Pair−wise Upper Bound − "Worst" pair of classes (2,3) (M=5)
Pair−wise Upper Bound − Pair of classes (1,3) (M=5)
(c)

Figure 1. Upper bound to error probability and true error probability *vs.* $1/\sigma^2$ (in dB) for: (a) the two zero-mean classes; (b) the two nonzero-mean classes; (c) multiple-classes.

$\mathbf{\Sigma}_4 = \mathbf{U}\text{diag}\,(0,0,0,0,1,1)\,\mathbf{U}^T$ where $\mathbf{U}$ is a randomly generated unitary matrix. Now, $r_{\mathbf{\Sigma}_1} = \text{rank}\,(\mathbf{\Sigma}_1) = 2$, $r_{\mathbf{\Sigma}_2} = \text{rank}\,(\mathbf{\Sigma}_2) = 3$, $r_{\mathbf{\Sigma}_3} = \text{rank}\,(\mathbf{\Sigma}_3) = 3$, $r_{\mathbf{\Sigma}_4} = \text{rank}\,(\mathbf{\Sigma}_4) = 2$, $r_{\mathbf{\Sigma}_{12}} = \text{rank}\,(\mathbf{\Sigma}_1 + \mathbf{\Sigma}_2) = 4$, $r_{\mathbf{\Sigma}_{13}} = \text{rank}\,(\mathbf{\Sigma}_1 + \mathbf{\Sigma}_3) = 5$, $r_{\mathbf{\Sigma}_{14}} = \text{rank}\,(\mathbf{\Sigma}_1 + \mathbf{\Sigma}_4) = 4$, $r_{\mathbf{\Sigma}_{23}} = \text{rank}\,(\mathbf{\Sigma}_2 + \mathbf{\Sigma}_3) = 4$, $r_{\mathbf{\Sigma}_{24}} = \text{rank}\,(\mathbf{\Sigma}_2 + \mathbf{\Sigma}_4) = 5$, $r_{\mathbf{\Sigma}_{34}} = \text{rank}\,(\mathbf{\Sigma}_3 + \mathbf{\Sigma}_4) = 4$ and $N = 6$. The sensing matrix is also generated randomly with i.i.d. zero-mean unit-variance Gaussian entries and normalized as previously noted. The pair of classes corresponding to the minimum pairwise diversity-order is $(2,3)$, offering a diversity-order that cannot exceed $-\frac{1}{2}\left(\frac{r_{\mathbf{\Sigma}_2}+r_{\mathbf{\Sigma}_3}}{2} - r_{\mathbf{\Sigma}_{23}}\right) = \frac{1}{2}$; all other pairs of classes offer a higher diversity-order.

Figure 1(c) confirms that the behavior of the upper bound to the misclassification probability is indeed dominated by the behavior of the pairwise upper bound associated with classes $(2,3)$. The Figure also confirms that the upper

bound to the misclassification probability approximates well the true misclassification probability: in particular, it is able to predict the presence of absence of the error floor as well as the diversity-order of the true error probability.

B. Designed Measurements

We now compare the performance of designed measurements to the performance of random measurements for two-class compressive classification problems with zero-mean and nonzero-mean classes and for zero-mean, multiple-class compressive classification problems.

For zero-mean two-class problems, we let $\boldsymbol{\mu}_1 = \boldsymbol{\mu}_2 = 0$ and for simplicity (and without loss of generality) $\mathbf{\Sigma}_1 = \mathrm{diag}\,(1,1,0)$ and $\mathbf{\Sigma}_2 = \mathrm{diag}\,(0,1,1)$. The realizations of the signals in class 1 live in the x_1–x_2 plane whereas the realizations of the signals in class 2 live in the x_2–x_3 plane; in addition, the dimension of the intersection of the sub-spaces associated with the signals in classes 1 and 2 is equal to one and the number of non-overlapping dimensions defined in (11) is equal to two. Realizations of the source are depicted in Figure 2(a).

The measurement matrix for this zero-mean two-class problem is constructed by taking the first M rows of the matrix:

$$\mathbf{\Phi}_0 = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix}, \tag{49}$$

according to the desired number of measurements M. Note that the first two rows of this matrix, which are consistent with the optimal design in Theorem 4, enable us to achieve the maximum diversity-order $d_{max} = \frac{1}{4}NO_{Dim} = \frac{1}{2}$ and the third row only provides for additional measurement gain.

For nonzero-mean two-class problems, we let $\boldsymbol{\mu}_1 = [0.328, 0.264, 0.114]^T$, $\boldsymbol{\mu}_2 = [1,1,1]^T$ and once again for simplicity $\mathbf{\Sigma}_1 = \mathbf{\Sigma}_2 = \mathrm{diag}\,(1,1,0)$. Note that the realizations of the signals in classes 1 and 2 are in two parallel affine spaces that differ by a translation corresponding to the vector $\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2$. Note also that the geometry is such that the condition in (32) in Theorem 5 is satisfied. Realizations of the source are depicted in Figure 2(b).

The measurement matrix for this nonzero-mean two-class problem is now constructed by taking the first M rows of the matrix:

$$\mathbf{\Phi}_0 = \begin{bmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix}, \tag{50}$$

according to the desired number of measurements. Note now that the first row of this matrix is in $\mathrm{Null}\,(\mathbf{\Sigma}_1 + \mathbf{\Sigma}_2)$, which in accordance with Theorem 5, provides for infinite diversity-order.

Figures 3(a) and 3(b) compare the performance of the random measurements to the designed measurements for the zero-mean two-class classification problem. One observes that $M = 3$ random measurements are necessary to eliminate the error floor in accordance with Theorem 2; however, a single $M = 1$ designed measurement is sufficient to drive the misclassification probability to zero in accordance with Theorem 4 due to the ability to focus on unique features exhibited by the classes. One also observes that it is possible to increase the diversity-order to $d_{max} = \frac{1}{4}NO_{Dim} = \frac{1}{2}$ by increasing the number of designed measurements from $M = 1$ to $M = 2$

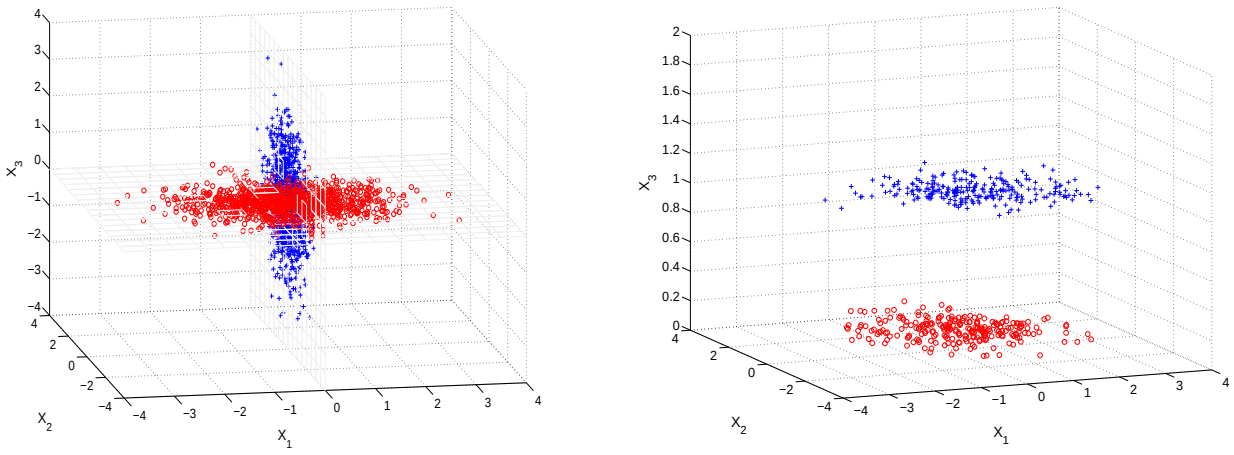

(a) (b)

Figure 2. Spatial representation of realizations of the source signals from classes 1 (in red circles) and 2 (in blue crosses) for a) the zero-mean two-class problem; and b) the nonzero-mean two-class problem.

and one additional measurement does not result in further increases in the diversity-order. Figures 4(a) and 4(b) instead compare the performance of random to designed measurements for the nonzero-mean two-class classification problem. In line with Theorems 3 and 5, $M = 3$ random measurements are necessary for the misclassification probability to decay exponentially as $\sigma^2 \to 0$ because with $M \leq r_{\mathbf{\Sigma}_{12}} = 2$ condition (21) can not be satisfied, but a single designed measurement is sufficient for the purpose. Overall, this behavior is also corroborated by the spatial representation of noiseless realizations of the projected signals portrayed in Figures 5 and 6: one can clearly see that fewer measurements are required in the designed case in relation to the random one to perfectly separate the classes.

For the multiple-class classification problem, we consider $L = 3$, $\boldsymbol{\mu}_1 = \boldsymbol{\mu}_2 = \boldsymbol{\mu}_3 = 0$, $\mathbf{\Sigma}_1 = \mathrm{diag}\,(1, 0, 0)$, $\mathbf{\Sigma}_2 = \mathrm{diag}\,(1, 1, 0)$ and $\mathbf{\Sigma}_3 = \mathrm{diag}\,(0, 1, 1)$.

We also construct a measurement matrix $\mathbf{\Phi}$ by taking the first M rows from the matrix

$$\mathbf{\Phi}_0 = \begin{bmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix}, \tag{51}$$

according to the number of measurements M. Note that this construction of $\mathbf{\Phi}$ follows the approach in Algorithm 1 for the measurement budget $M = 3$, as each row represents a measurement that achieves diversity-order $1/4$ for a given pair of classes. The advantage of using designed measurements in *lieu* of random ones is also apparent for multiple-class problems. In the designed case one needs only two measurements to drive the misclassification probability to zero as the noise tends to zero with maximal diversity-order. In contrast, with random measurements, one needs at least three random measurements to be able to eliminate the error floor.

Finally, we also notice that the upper bound to the misclassification probability captures well the behavior of the

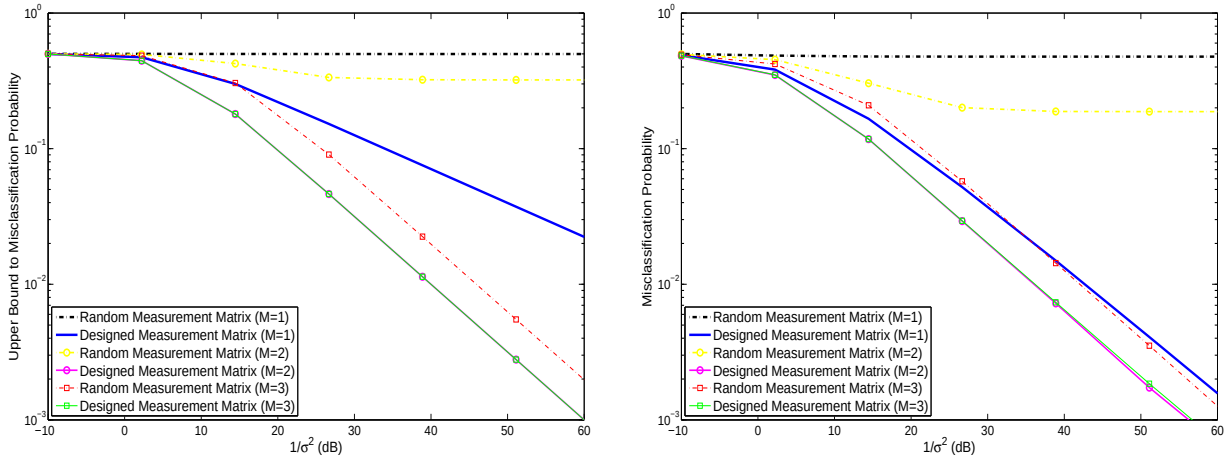

(a) (b)

Figure 3. Upper bound to the probability of misclassification (a) and true probability of misclassification (b) vs. $1/\sigma^2$ (in dB) for random and designed measurements (two zero-mean classes).

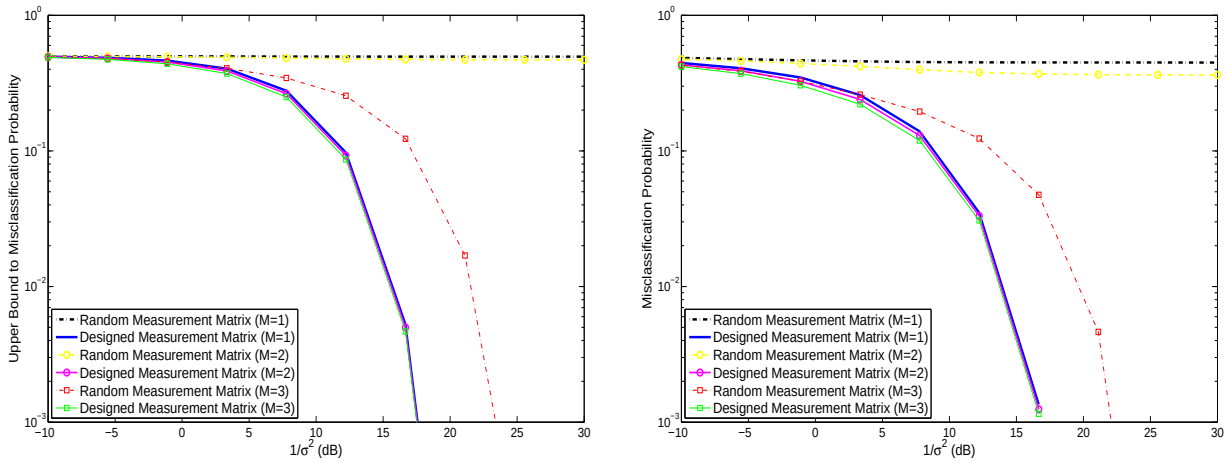

(a) (b)

Figure 4. Upper bound to the probability of misclassification (a) and true probability of misclassification (b) vs. $1/\sigma^2$ (in dB) for random and designed measurements (two nonzero-mean classes).

error probability: in particular, it captures well the presence of absence of an error floor, the true diversity-order, increases in diversity-order and increases in the measurement gain. In any case, the diversity-order associated with the upper bound to the misclassification probability always lower bounds the diversity-order associated with the true misclassification probability, thereby establishing performance assurances.

VIII. Concluding Remarks

This paper studies fundamental limits in the compressive classification of a mixture of Gaussians by using performance characterizations that are the duals of performance characterizations in multiple-antenna communications systems. In particular, by considering the diversity-order and measurement gain associated with a Bhattacharrya

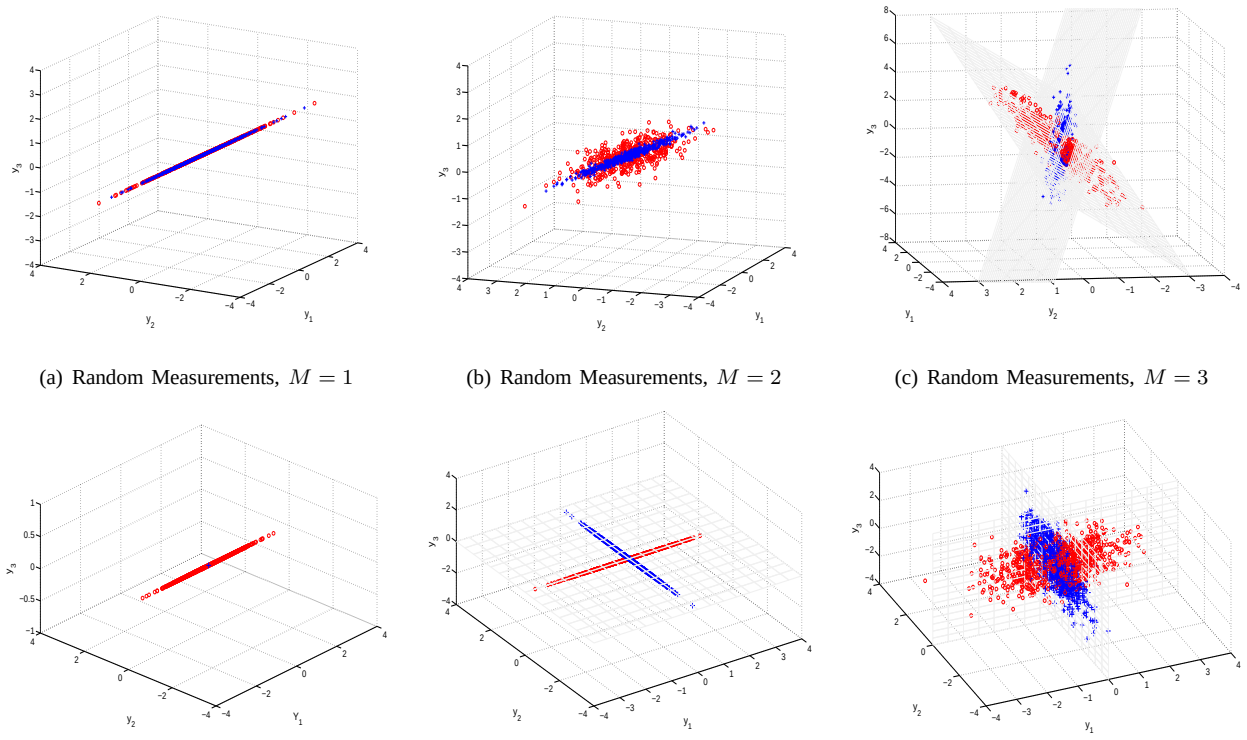

(a) Random Measurements, $M = 1$ (b) Random Measurements, $M = 2$ (c) Random Measurements, $M = 3$

(d) Designed Measurements, $M = 1$ (e) Designed Measurements, $M = 2$ (f) Designed Measurements, $M = 3$

Figure 5. Spatial representation of realizations of noiseless projected source signals from classes 1 (in red circles) and 2 (in blue crosses) for zero-mean classes: (a) Random Measurements, $M = 1$; (b) Random Measurements, $M = 2$; (c) Random Measurements, $M = 3$; (d) Designed Measurements, $M = 1$; (e) Designed Measurements, $M = 2$; (f) Designed Measurements, $M = 3$.

based upper bound to the misclassification probability of the compressive classification problem, which act as the counterparts to the diversity-order and coding gain associated with upper bounds to the error probability of a multiple-antenna communications problem, it has been possible to provide more refined performance characterizations that capture well not only the presence or absence of misclassification error floors (a phase transition) but also increases in the diversity-order and increases in the measurement gain associated with the true misclassification probability of our compressive classification problem.

The proposed characterizations have been used to study the performance of two-class classification problems with zero-mean Gaussians, two-class classification problems with nonzero-mean Gaussians, and multiple-class Gaussian classification problems, both in the presence of random Gaussian i.i.d. measurements and (diversity-order) optimized measurements. One of the hallmarks of the proposed characterizations is the ability to link the concepts of diversity gain and measurement gain with certain fundamental geometrical quantities associated with the measurement and the source models. For example, it has been shown that the ultimate diversity-order in a two-class compressive classification problem with zero-mean Gaussians and, subject to some mild conditions, with nonzero-mean Gaussians, is dictated by the so-called number of non-overlapping dimensions, a quantity that can be

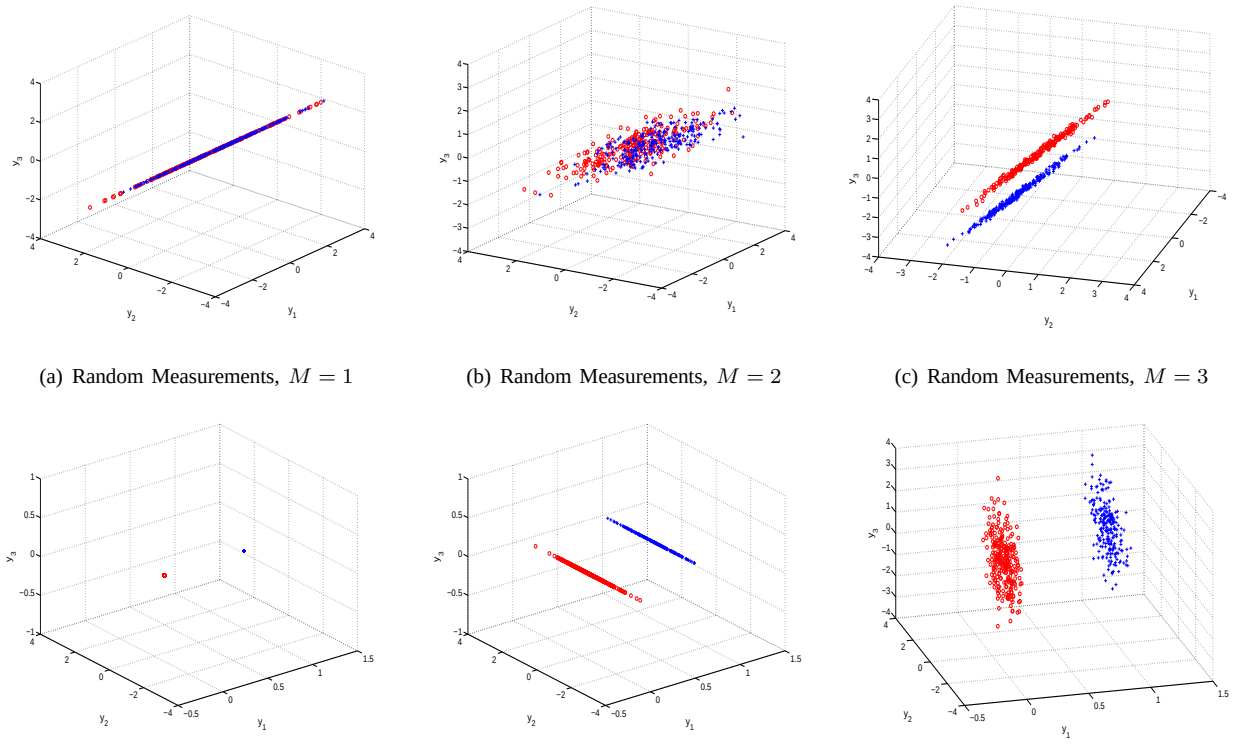

(a) Random Measurements, $M = 1$ (b) Random Measurements, $M = 2$ (c) Random Measurements, $M = 3$

(d) Designed Measurements, $M = 1$ (e) Designed Measurements, $M = 2$ (f) Designed Measurements, $M = 3$

Figure 6. Spatial representation of realizations of noiseless projected source signals from classes 1 (in red circles) and 2 (in blue crosses) for nonzero-mean classes: (a) Random Measurements, $M = 1$; (b) Random Measurements, $M = 2$; (c) Random Measurements, $M = 3$; (d) Designed Measurements, $M = 1$; (e) Designed Measurements, $M = 2$; (f) Designed Measurements, $M = 3$.

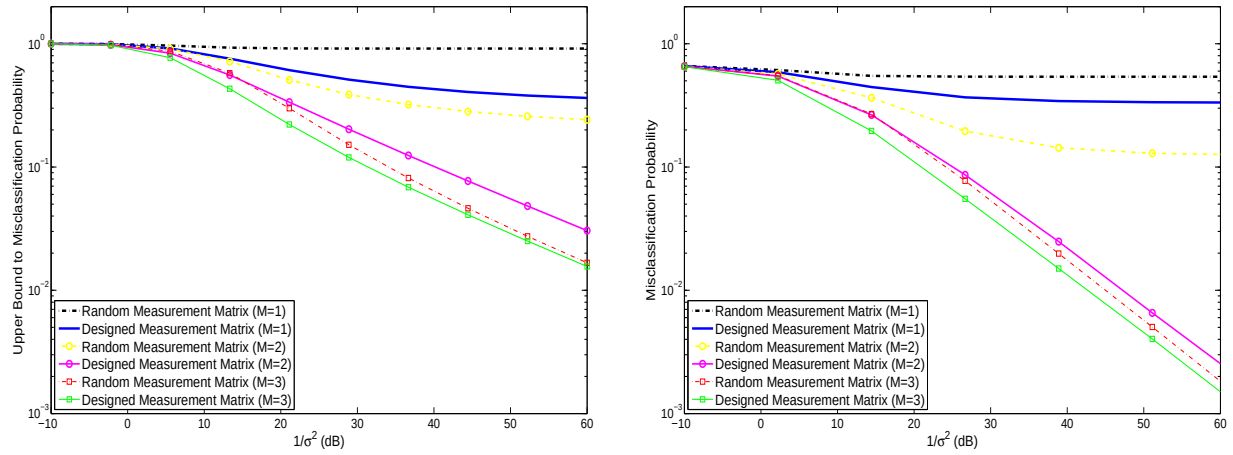

(a) (b)

Figure 7. Upper bound to the probability of misclassification (a) and true probability of misclassification (b) vs $1/\sigma^2$ (in dB) for random and designed measurements (multiple classes).

interpreted as the number of unique features associated with two-class classification problems. One then understands that measurement is a means to probe such unique features and that designed measurements provide a quicker route to probe such features in relation to the standard random ones.

Finally, it is also relevant to remark that one possible ramification of the asymptotic performance characterizations and results concerns dictionary learning and design for compressive classification problems. For example, if we are allowed to jointly optimize $\mathbf{\Phi}$, $\mathbf{\Sigma}_1$ and $\mathbf{\Sigma}_2$ in a two-class problem with zero-mean Gaussians, one can easily show that the diversity-maximizing design is such that $r_{12} = \mathrm{rank}\left(\mathbf{\Phi}\left(\mathbf{\Sigma}_1+\mathbf{\Sigma}_2\right)\mathbf{\Phi}^T\right) = N$ and $r_1 + r_2 = \mathrm{rank}\left(\mathbf{\Phi}\mathbf{\Sigma}_1\mathbf{\Phi}^T\right) + \mathrm{rank}\left(\mathbf{\Phi}\mathbf{\Sigma}_2\mathbf{\Phi}^T\right) = r_{12} = N$. This design procedure is in fact reminiscent of recent state-of-the-art methods associated with learning transformations for subspace clustering and classification, that seek to construct a linear transformation on subspaces using matrix rank via its convex surrogate nuclear norm: concretely, the goal is to learn a linear transformation that enforces a low-rank structure for data from the same subspace and that also enforces a high-rank structure for data from different subspaces [82].

APPENDIX A
PROOF OF THEOREM 1

Consider the eigenvalue decomposition of the following matrices $\mathbf{\Phi}\mathbf{\Sigma}_1\mathbf{\Phi}^T = \mathbf{U}_1\mathbf{\Lambda}_1\mathbf{U}_1^T$, $\mathbf{\Phi}\mathbf{\Sigma}_2\mathbf{\Phi}^T = \mathbf{U}_2\mathbf{\Lambda}_2\mathbf{U}_2^T$, $\mathbf{\Phi}\left(\mathbf{\Sigma}_1+\mathbf{\Sigma}_2\right)\mathbf{\Phi}^T = \mathbf{U}_{12}\mathbf{\Lambda}_{12}\mathbf{U}_{12}^T$, where $\mathbf{U}_1, \mathbf{U}_2, \mathbf{U}_{12} \in \mathbb{R}^{M\times M}$ are orthogonal matrices and $\mathbf{\Lambda}_1$, $\mathbf{\Lambda}_2$ and $\mathbf{\Lambda}_{12}$ are positive semidefinite diagonal matrices such that $\mathbf{\Lambda}_1 = \mathrm{diag}\left(\lambda_{1_1}, \cdots, \lambda_{1_{r_1}}, 0, \cdots, 0\right)$, $\mathbf{\Lambda}_2 = \mathrm{diag}\left(\lambda_{2_1}, \cdots, \lambda_{2_{r_2}}, 0, \cdots, 0\right)$, $\mathbf{\Lambda}_{12} = \mathrm{diag}\left(\lambda_{12_1}, \cdots, \lambda_{12_{r_{12}}}, 0, \cdots, 0\right)$; and $r_1 = \mathrm{rank}\left(\mathbf{\Phi}\mathbf{\Sigma}_1\mathbf{\Phi}^T\right)$, $r_2 = \mathrm{rank}\left(\mathbf{\Phi}\mathbf{\Sigma}_2\mathbf{\Phi}^T\right)$ and $r_{12} = \mathrm{rank}\left(\mathbf{\Phi}\left(\mathbf{\Sigma}_1+\mathbf{\Sigma}_2\right)\mathbf{\Phi}^T\right)$.

Therefore, we can re-express the upper bound to the misclassification error probability in (14) as follows:

$$P_{err}^{UB} = \sqrt{P_1P_2}\; e^{\left[-\frac{1}{2}\log\frac{\det\left(\frac{\mathbf{\Phi}\left(\mathbf{\Sigma}_1+\mathbf{\Sigma}_2\right)\mathbf{\Phi}^T+2\sigma^2\mathbf{I}}{2}\right)}{\sqrt{\det\left(\mathbf{\Phi}\mathbf{\Sigma}_1\mathbf{\Phi}^T+\sigma^2\mathbf{I}\right)\det\left(\mathbf{\Phi}\mathbf{\Sigma}_2\mathbf{\Phi}^T+\sigma^2\mathbf{I}\right)}}\right]}$$
$$= \sqrt{P_1P_2}\; e^{\left[-\frac{1}{2}\log\left[2^{-r_{12}}\left(\sigma^2\right)^{\frac{r_1+r_2}{2}-r_{12}}\frac{\prod_{i=1}^{r_{12}}\left(\lambda_{12_i}+2\sigma^2\right)}{\sqrt{\prod_{i=1}^{r_1}\left(\lambda_{1_i}+\sigma^2\right)\prod_{i=1}^{r_2}\left(\lambda_{2_i}+\sigma^2\right)}}\right]\right]}. \quad (52)$$

The asymptotic characterization of the behavior of the upper bound to the probability of misclassification follows immediately from (3) and (4) together with (52). In particular,

- If $\frac{r_1+r_2}{2} = r_{12}$ then, $\lim_{\sigma^2\to 0} P_{err}^{UB} = \sqrt{P_1P_2}\left[2^{-r_{12}}\frac{v_{12}}{\sqrt{v_1v_2}}\right]^{-\frac{1}{2}} \neq 0$,

- If $\frac{r_1+r_2}{2} < r_{12}$ then, $\lim_{\sigma^2\to 0} P_{err}^{UB} = 0$, and by using (52) in (3) and in (4):

$$d = -\frac{1}{2}\left(\frac{r_1+r_2}{2} - r_{12}\right) \quad (53)$$

and

$$g_m = \left[2^{\frac{r_{12}}{2}}\sqrt{P_1P_2}\left[\frac{v_{12}}{\sqrt{v_1v_2}}\right]^{-\frac{1}{2}}\right]^{-\frac{1}{d}}. \quad (54)$$

APPENDIX B
PROOF OF THEOREM 2

Consider once again the eigenvalue decompositions in Appendix A. Consider also that $N \geq r_{\mathbf{\Sigma}_{12}} \geq \max\left(r_{\mathbf{\Sigma}_1}, r_{\mathbf{\Sigma}_2}\right)$ and, with probability 1, $r_1 = \min\left(M, r_{\mathbf{\Sigma}_1}\right)$, $r_2 = \min\left(M, r_{\mathbf{\Sigma}_2}\right)$ and $r_{12} = \min\left(M, r_{\mathbf{\Sigma}_{12}}\right)$. In addition, assume, without any loss of generality, that $r_{\mathbf{\Sigma}_1} \leq r_{\mathbf{\Sigma}_2}$. Therefore, we can re-express the upper bound to the misclassification error probability in (14) as follows:

$$P_{err}^{UB} = \sqrt{P_1 P_2}\; e^{\left[-\frac{1}{2}\log \frac{\det\left(\frac{\mathbf{\Phi}(\mathbf{\Sigma}_1+\mathbf{\Sigma}_2)\mathbf{\Phi}^T+2\sigma^2\mathbf{I}}{2}\right)}{\sqrt{\det\left(\mathbf{\Phi}\mathbf{\Sigma}_1\mathbf{\Phi}^T+\sigma^2\mathbf{I}\right)\det\left(\mathbf{\Phi}\mathbf{\Sigma}_2\mathbf{\Phi}^T+\sigma^2\mathbf{I}\right)}}\right]}$$
$$= \sqrt{P_1 P_2}\; e^{\left[-\frac{1}{2}\log\left[2^{-\min\left(M,r_{\mathbf{\Sigma}_{12}}\right)}\left(\sigma^2\right)^{\frac{\min\left(M,r_{\mathbf{\Sigma}_1}\right)+\min\left(M,r_{\mathbf{\Sigma}_2}\right)}{2}-\min\left(M,r_{\mathbf{\Sigma}_{12}}\right)}\frac{\prod_{i=1}^{\min\left(M,r_{\mathbf{\Sigma}_{12}}\right)}\left(\lambda_{12_i}+2\sigma^2\right)}{\sqrt{\prod_{i=1}^{r_1}\left(\lambda_{1_i}+\sigma^2\right)\prod_{i=1}^{r_2}\left(\lambda_{2_i}+\sigma^2\right)}}\right]\right]}. \quad (55)$$

The asymptotic characterization of the behavior of the upper bound to the probability of misclassification follows, once again, immediately from (3) and (4) together with (55). In particular,

- If $\frac{r_{\mathbf{\Sigma}_1}+r_{\mathbf{\Sigma}_2}}{2} = r_{\mathbf{\Sigma}_{12}}$ then, $\lim_{\sigma^2\to 0} P_{err}^{UB} = \sqrt{P_1 P_2}\left[2^{-\min\left(M,r_{\mathbf{\Sigma}_{12}}\right)}\frac{v_{12}}{\sqrt{v_1 v_2}}\right]^{-\frac{1}{2}} \neq 0$,
- If $\frac{r_{\mathbf{\Sigma}_1}+r_{\mathbf{\Sigma}_2}}{2} < r_{\mathbf{\Sigma}_{12}}$ and $M \leq r_{\mathbf{\Sigma}_1} \leq r_{\mathbf{\Sigma}_2} \leq r_{\mathbf{\Sigma}_{12}}$ then, $\lim_{\sigma^2\to 0} P_{err}^{UB} = \left[2^{-M}\frac{v_{12}}{\sqrt{v_1 v_2}}\right]^{-\frac{1}{2}} \neq 0$,
- otherwise, $\lim_{\sigma^2\to 0} P_{err}^{UB} = 0$, and, by using in (55) in (4) and in (3), we can write the measurement gain as:

$$g_m = \left[2^{\frac{\min\left(M,r_{\mathbf{\Sigma}_{12}}\right)}{2}}\sqrt{P_1 P_2}\left[\frac{v_{12}}{\sqrt{v_1 v_2}}\right]^{-\frac{1}{2}}\right]^{-\frac{1}{d}} \quad (56)$$

and the diversity-order as:

$$d = \begin{cases} -\frac{1}{2}\left(\frac{r_{\mathbf{\Sigma}_1}-M}{2}\right) & , \text{ if } r_{\mathbf{\Sigma}_1} < M \leq r_{\mathbf{\Sigma}_2} \leq r_{\mathbf{\Sigma}_{12}} \\ -\frac{1}{2}\left(\frac{r_{\mathbf{\Sigma}_1}+r_{\mathbf{\Sigma}_2}}{2}-M\right) & , \text{ if } r_{\mathbf{\Sigma}_1} \leq r_{\mathbf{\Sigma}_2} < M < r_{\mathbf{\Sigma}_{12}} \\ -\frac{1}{2}\left(\frac{r_{\mathbf{\Sigma}_1}+r_{\mathbf{\Sigma}_2}}{2}-r_{\mathbf{\Sigma}_{12}}\right) & , \text{ if } r_{\mathbf{\Sigma}_1} \leq r_{\mathbf{\Sigma}_2} \leq r_{\mathbf{\Sigma}_{12}} \leq M \end{cases}. \quad (57)$$

APPENDIX C
PROOF OF THEOREM 3

Consider the upper bound to the classification error probability in (14) and (15). We write the exponent K_{12} as follows:

$$K_{12} = T_1 + T_2 \quad (58)$$

where

$$T_1 = \frac{1}{8}\left[\mathbf{\Phi}\left(\boldsymbol{\mu}_1-\boldsymbol{\mu}_2\right)\right]^T\left[\frac{\mathbf{\Phi}\left(\mathbf{\Sigma}_1+\mathbf{\Sigma}_2\right)\mathbf{\Phi}^T+2\sigma^2\mathbf{I}}{2}\right]^{-1}\left[\mathbf{\Phi}\left(\boldsymbol{\mu}_1-\boldsymbol{\mu}_2\right)\right] \quad (59)$$

and

$$T_2 = \frac{1}{2}\log\frac{\det\left(\frac{\mathbf{\Phi}(\mathbf{\Sigma}_1+\mathbf{\Sigma}_2)\mathbf{\Phi}^T+2\sigma^2\mathbf{I}}{2}\right)}{\sqrt{\det\left(\mathbf{\Phi}\mathbf{\Sigma}_1\mathbf{\Phi}^T+\sigma^2\mathbf{I}\right)\det\left(\mathbf{\Phi}\mathbf{\Sigma}_2\mathbf{\Phi}^T+\sigma^2\mathbf{I}\right)}}. \tag{60}$$

Let us now define $\mathbf{M}_{12} = \mathbf{\Phi}(\boldsymbol{\mu}_1-\boldsymbol{\mu}_2)(\boldsymbol{\mu}_1-\boldsymbol{\mu}_2)^T\mathbf{\Phi}^T$. Then, by recalling the eigenvalue decomposition of the matrix $\mathbf{\Phi}\left(\mathbf{\Sigma}_1+\mathbf{\Sigma}_2\right)\mathbf{\Phi}^T = \mathbf{U}_{12}\mathbf{\Lambda}_{12}\mathbf{U}_{12}^T$ in Appendix A, we can also express T_1 as:

$$T_1 = \frac{1}{4}\mathrm{tr}\left(\mathbf{M}_{12}\left(\mathbf{\Phi}(\mathbf{\Sigma}_1+\mathbf{\Sigma}_2)\mathbf{\Phi}^T+2\sigma^2\mathbf{I}\right)^{-1}\right) \tag{61}$$

$$= \frac{1}{4}\mathrm{tr}\left(\mathbf{M}_{12}\mathbf{U}_{12}\left(\mathbf{\Lambda}_{12}+2\sigma^2\mathbf{I}\right)^{-1}\mathbf{U}_{12}^T\right) \tag{62}$$

$$= \frac{1}{4}\sum_{i=1}^{r_{12}}\frac{1}{\lambda_{12_i}+2\sigma^2}\mathbf{u}_{12,i}^T\mathbf{M}_{12}\mathbf{u}_{12,i}+\frac{1}{8\sigma^2}\sum_{i=r_{12}+1}^{M}\mathbf{u}_{12,i}^T\mathbf{M}_{12}\mathbf{u}_{12,i}, \tag{63}$$

where the vector $\mathbf{u}_{12,i}$ corresponds to the i-th column of the matrix $\mathbf{U}_{12}$.

Therefore, we can also re-write the exponent K_{12} as follows:

$$K_{12} = \frac{1}{4}\sum_{i=1}^{r_{12}}\frac{1}{\lambda_{12_i}+2\sigma^2}\mathbf{u}_{12,i}^T\mathbf{M}_{12}\mathbf{u}_{12,i}+\frac{1}{8\sigma^2}\sum_{i=r_{12}+1}^{M}\mathbf{u}_{12,i}^T\mathbf{M}_{12}\mathbf{u}_{12,i}+T_2, \tag{64}$$

We are now able to extend the characterization of the asymptotic behavior of the Batthacharyya upper bound from the case of zero-mean to nonzero-mean classes by analyzing further (64). In particular,

- If $\boldsymbol{\mu}_1 = \boldsymbol{\mu}_2$ then the asymptotic behavior of the upper bound to the misclassification probability is identical in nonzero-mean and zero-mean cases;
- If $\boldsymbol{\mu}_1 \neq \boldsymbol{\mu}_2$ and $M \leq r_{\mathbf{\Sigma}_{12}}$, then the matrix $\mathbf{\Phi}\left(\mathbf{\Sigma}_1+\mathbf{\Sigma}_2\right)\mathbf{\Phi}^T$ is full rank (i.e. $r_{12} = M$) and, therefore, condition (21) is not verified. In such case the second term in (64) is equal to zero and the diversity-order associated with nonzero-mean classes corresponds to that for zero-mean classes unveiled in Theorem 2. In contrast, the measurement gain for nonzero-mean classes is higher than that for zero-mean classes in view of the first term in (64). In fact, it is immediate to express the measurement gain for nonzero-mean classes g_m^{NZM} in terms of the measurement gain for zero-mean classes g_m^{ZM} in (19) as $g_m^{\mathrm{NZM}} = a \cdot g_m^{\mathrm{ZM}}$ where

$$a = \exp\left(\frac{1}{4d}\sum_{i=1}^{M}\frac{1}{\lambda_{12_i}}\mathbf{u}_{12,i}^T\mathbf{M}_{12}\mathbf{u}_{12,i}\right) > 1. \tag{65}$$

- If $\boldsymbol{\mu}_1 \neq \boldsymbol{\mu}_2$ and $M > r_{\mathbf{\Sigma}_{12}}$, then the second term in (64) is given by

$$\frac{1}{8\sigma^2}\sum_{i=r_{\mathbf{\Sigma}_{12}}+1}^{M}\mathbf{u}_{12,i}^T\mathbf{M}_{12}\mathbf{u}_{12,i}, \tag{66}$$

  where the vectors $\mathbf{u}_{12,r_{\mathbf{\Sigma}_{12}}+1},\ldots,\mathbf{u}_{12,M}$ are the eigenvectors of the matrix $\mathbf{\Phi}(\mathbf{\Sigma}_1+\mathbf{\Sigma}_2)\mathbf{\Phi}^T$ corresponding to the zero eigenvalues, which form an orthonormal basis of the null space of that matrix. In this case, since the matrices $\mathbf{\Phi}(\mathbf{\Sigma}_1+\mathbf{\Sigma}_2)\mathbf{\Phi}^T$ and $\mathbf{M}_{12}$ are positive semi-definite, by using the fundamental theorem of algebra, we can conclude that (66) is equal to zero if and only if

$$\mathrm{Null}(\mathbf{\Phi}(\mathbf{\Sigma}_1+\mathbf{\Sigma}_2)\mathbf{\Phi}^T) \subseteq \mathrm{Null}(\mathbf{M}_{12}) \Leftrightarrow \mathrm{im}(\mathbf{M}_{12}) \subseteq \mathrm{im}(\mathbf{\Phi}(\mathbf{\Sigma}_1+\mathbf{\Sigma}_2)\mathbf{\Phi}^T) \Leftrightarrow \mathbf{\Phi}(\boldsymbol{\mu}_1-\boldsymbol{\mu}_2) \in \mathrm{im}(\mathbf{\Phi}(\mathbf{\Sigma}_1+\mathbf{\Sigma}_2)\mathbf{\Phi}^T \tag{67}$$

Therefore, if we do not satisfy (67) then (66) is strictly greater than zero and the upper bound to the misclassification probability decays exponentially with $1/\sigma^2$ as $\sigma^2 \to 0$.

APPENDIX D
PROOF OF THEOREM 4

Recall that the diversity-order is given by:

$$d = -\frac{1}{2}\left(\frac{r_1 + r_2}{2} - r_{12}\right) = \frac{1}{4}\left(2r_{12} - r_1 - r_2\right). \tag{68}$$

We establish the designs by pursuing a two-step approach where we first determine an upper bound to the maximum diversity-order and we then determine a design that achieves such a maximum diversity-order. We consider the cases where $M \geq NO_{Dim}$ and $M < NO_{Dim}$ separately.

A. Case Where $M \geq NO_{Dim}$

The maximum diversity-order that we can achieve with any measurement matrix when $M \geq NO_{Dim}$ is given by:

$$d \leq \frac{1}{4} NO_{Dim}. \tag{69}$$

This upper bound can be proven by showing that:

$$2r_{12} - r_1 - r_2 \leq 2r_{\mathbf{\Sigma}_{12}} - r_{\mathbf{\Sigma}_1} - r_{\mathbf{\Sigma}_2} = NO_{Dim}, \tag{70}$$

or, instead,

$$r_{\mathbf{\Sigma}_{12}} - r_{12} \geq r_{\mathbf{\Sigma}_1} - r_1 \wedge r_{\mathbf{\Sigma}_{12}} - r_{12} \geq r_{\mathbf{\Sigma}_2} - r_2, \tag{71}$$

since (71) implies (70). Consider the generalized eigenvalue decomposition of the positive semidefinite matrices $\mathbf{\Sigma}_1$ and $\mathbf{\Sigma}_2$ given by [83, Theorem 8.7.1], namely, $\mathbf{\Sigma}_1 = \mathbf{X}^{-T}\mathbf{D}_1\mathbf{X}^{-1} = \mathbf{X}^{-T}\ \mathrm{diag}\left(d_{1_1}, \ldots, d_{1_N}\right)\mathbf{X}^{-1}$ with $d_{1_i} \geq 0, i = 1, \ldots, N$ and $\mathbf{\Sigma}_2 = \mathbf{X}^{-T}\mathbf{D}_2\mathbf{X}^{-1} = \mathbf{X}^{-T}\ \mathrm{diag}\left(d_{2_1}, \ldots, d_{2_N}\right)\mathbf{X}^{-1}$ with $d_{2_i} \geq 0, i = 1, \ldots, N$, where $\mathbf{X}$ is a non-singular matrix.

Now, we can write $r_{12} = \mathrm{rank}\left(\mathbf{\Phi}\mathbf{X}^{-T}\left(\mathbf{D}_1 + \mathbf{D}_2\right)\mathbf{X}^{-1}\mathbf{\Phi}^{-T}\right) = \mathrm{rank}\left(\tilde{\mathbf{\Phi}}\left(\mathbf{D}_1 + \mathbf{D}_2\right)^{\frac{1}{2}}\right)$ and likewise, $r_1 = \mathrm{rank}\left(\mathbf{\Phi}\mathbf{X}^{-T}\left(\mathbf{D}_1\right)\mathbf{X}^{-1}\mathbf{\Phi}^{-T}\right) = \mathrm{rank}\left(\tilde{\mathbf{\Phi}}\left(\mathbf{D}_1\right)^{\frac{1}{2}}\right)$ and $r_2 = \mathrm{rank}\left(\mathbf{\Phi}\mathbf{X}^{-T}\left(\mathbf{D}_2\right)\mathbf{X}^{-1}\mathbf{\Phi}^{-T}\right) = \mathrm{rank}\left(\tilde{\mathbf{\Phi}}\left(\mathbf{D}_2\right)^{\frac{1}{2}}\right)$, where $\tilde{\mathbf{\Phi}} = \mathbf{\Phi}\mathbf{X}^{-T}$.

On the other hand, the ranks of the input covariance matrices can be expressed as $r_{\mathbf{\Sigma}_{12}} = \mathrm{rank}\left(\mathbf{X}^{-T}\left(\mathbf{D}_1 + \mathbf{D}_2\right)\mathbf{X}^{-1}\right) = \mathrm{rank}\left(\mathbf{X}^{-T}\left(\mathbf{D}_1 + \mathbf{D}_2\right)^{\frac{1}{2}}\right) = \mathrm{rank}\left(\left(\mathbf{D}_1 + \mathbf{D}_2\right)^{\frac{1}{2}}\right)$ and $r_{\mathbf{\Sigma}_1} = \mathrm{rank}\left(\left(\mathbf{D}_1\right)^{\frac{1}{2}}\right)$ and $r_{\mathbf{\Sigma}_2} = \mathrm{rank}\left(\left(\mathbf{D}_2\right)^{\frac{1}{2}}\right)$.

Let us now define the cardinalities of the following sets: $k_c = |\{i : d_{1_i} > 0 \wedge d_{2_i} > 0\}|$, $k_1 = |\{i : d_{1_i} > 0\}|$ and $k_2 = |\{i : d_{2_i} > 0\}|$. Then, it becomes evident that, $r_{\mathbf{\Sigma}_{12}} - r_{\mathbf{\Sigma}_1} = k_1 + k_2 - k_c - k_1 = k_2 - k_c$ and $r_{\mathbf{\Sigma}_{12}} - r_{\mathbf{\Sigma}_2} = k_1 + k_2 - k_c - k_2 = k_1 - k_c$, and, in view of the possible dependence between columns of $\tilde{\mathbf{\Phi}}$, $r_{12} - r_1 \leq k_2 - k_c$ and $r_{12} - r_2 \leq k_1 - k_c$, thus concluding the proof for the upper bound on the diversity-order.

This upper bound – in view of (70) – can be achieved by a projections matrix design that satisfies

$$NO_{Dim} = 2r_{12} - r_1 - r_2 = 2r_{\mathbf{\Sigma}_{12}} - r_{\mathbf{\Sigma}_1} - r_{\mathbf{\Sigma}_2} = n_{\mathbf{\Sigma}_1} + n_{\mathbf{\Sigma}_2}, \tag{72}$$

where we have used the fact that $r_{\mathbf{\Sigma}_1} = N - \dim\left(\mathrm{Null}\left(\mathbf{\Sigma}_1\right)\right) = N - n_{12} - n_{\mathbf{\Sigma}_1}$, $r_{\mathbf{\Sigma}_2} = N - \dim\left(\mathrm{Null}\left(\mathbf{\Sigma}_2\right)\right) = N - n_{12} - n_{\mathbf{\Sigma}_2}$ and $r_{\mathbf{\Sigma}_{12}} = N - \dim\left(\mathrm{Null}\left(\mathbf{\Sigma}_1\right) \bigcap \mathrm{Null}\left(\mathbf{\Sigma}_2\right)\right) = N - n_{12}$, where n_{12}, $n_{12}+n_{\mathbf{\Sigma}_1}$ and $n_{12}+n_{\mathbf{\Sigma}_2}$ are the dimensions of the sub-spaces $\mathrm{Null}\left(\mathbf{\Sigma}_1\right) \bigcap \mathrm{Null}\left(\mathbf{\Sigma}_2\right)$, $\mathrm{Null}\left(\mathbf{\Sigma}_1\right)$ and $\mathrm{Null}\left(\mathbf{\Sigma}_2\right)$, respectively. Note that, in order to guarantee (72), the two conditions in (71) have to hold with equality, thus implying that $r_2 \geq 2r_{12} - r_1 = r_{\mathbf{\Sigma}_{12}} - r_{\mathbf{\Sigma}_1}$ and $r_1 \geq 2r_{12} - r_2 = r_{\mathbf{\Sigma}_{12}} - r_{\mathbf{\Sigma}_2}$.

A possible measurement matrix construction that achieves the maximum diversity-order is

$$\mathbf{\Phi} = \left[\boldsymbol{v}_1, \boldsymbol{v}_2, \ldots, \boldsymbol{v}_{n_{\mathbf{\Sigma}_1}}, \boldsymbol{w}_1, \boldsymbol{w}_2, \ldots, \boldsymbol{w}_{n_{\mathbf{\Sigma}_2}}\right]^T, \tag{73}$$

where the set of vectors $\left[\boldsymbol{u}_1, \ldots, \boldsymbol{u}_{n_{12}}\right]$, $\left[\boldsymbol{u}_1, \ldots, \boldsymbol{u}_{n_{12}}, \boldsymbol{v}_1, \ldots, \boldsymbol{v}_{n_{\mathbf{\Sigma}_1}}\right]$, $\left[\boldsymbol{u}_1, \ldots, \boldsymbol{u}_{n_{12}}, \boldsymbol{w}_1, \ldots, \boldsymbol{w}_{n_{\mathbf{\Sigma}_2}}\right]$, $\boldsymbol{u}_i, \boldsymbol{v}_i, \boldsymbol{w}_i \in \mathbb{R}^N$, constitute an orthonormal basis of the linear spaces $\mathrm{Null}\left(\mathbf{\Sigma}_1\right) \bigcap \mathrm{Null}\left(\mathbf{\Sigma}_2\right)$, $\mathrm{Null}\left(\mathbf{\Sigma}_1\right)$ and $\mathrm{Null}\left(\mathbf{\Sigma}_2\right)$, respectively. This can be verified by writing

$$\mathbf{\Phi}\mathbf{\Sigma}_1\mathbf{\Phi}^T = \left[\begin{array}{c|c} \mathbf{0} & \mathbf{0} \\ \hline \mathbf{0} & \mathbf{Q} \end{array}\right] \quad , \quad \mathbf{Q} = \left[\boldsymbol{w}_1 \quad \boldsymbol{w}_2 \quad \cdots \quad \boldsymbol{w}_{n_{\mathbf{\Sigma}_2}}\right]^T \mathbf{\Sigma}_1 \left[\boldsymbol{w}_1 \quad \boldsymbol{w}_2 \quad \cdots \quad \boldsymbol{w}_{n_{\mathbf{\Sigma}_2}}\right] \tag{74}$$

where $r_1 = \mathrm{rank}\left(\mathbf{\Phi}\mathbf{\Sigma}_1\mathbf{\Phi}^T\right) = \mathrm{rank}\left(\mathbf{Q}\right)$.

Now, note that the matrix $\mathbf{Q}$ is the Gram matrix of the set of vectors $\boldsymbol{q}_i = \mathbf{\Sigma}_1^{\frac{1}{2}}\boldsymbol{w}_i$, $i = 1, \ldots, n_{\mathbf{\Sigma}_2}$, and, therefore, $r_1 = \mathrm{rank}\left(\mathbf{Q}\right) = n_{\mathbf{\Sigma}_2}$ if and only if the vectors $\boldsymbol{q}_i$, $i = 1, \ldots, n_{\mathbf{\Sigma}_2}$, are linearly independent.

Assume by contradiction that the vectors $\boldsymbol{q}_i$ are linearly dependent. Then, there exists a set of $n_{\mathbf{\Sigma}_2}$ scalars α_i (with $\alpha_i \neq 0$ for at least one index i) such that $\mathbf{\Sigma}_1^{\frac{1}{2}} \sum_i \alpha_i \boldsymbol{w}_i = \mathbf{0}$. It is known that $\sum_i \alpha_i \boldsymbol{w}_i \neq \mathbf{0}$ because $\boldsymbol{w}_i$ are linearly independent by construction. Therefore, the linearly dependence among the vectors $\boldsymbol{q}_i$ implies that

$$\sum_i \alpha_i \boldsymbol{w}_i \in \mathrm{Null}\left(\mathbf{\Sigma}_1^{\frac{1}{2}}\right) \text{ or } \sum_i \alpha_i \boldsymbol{w}_i \in \mathrm{Null}\left(\mathbf{\Sigma}_1\right) \tag{75}$$

which is false since, by construction, $\sum_i \alpha_i \boldsymbol{w}_i \in \mathrm{Null}\left(\mathbf{\Sigma}_2\right)$ and $\sum_i \alpha_i \boldsymbol{w}_i \notin \mathrm{Null}\left(\mathbf{\Sigma}_1\right) \bigcap \mathrm{Null}\left(\mathbf{\Sigma}_2\right)$. Therefore, we can establish that $r_1 = \mathrm{rank}\left(\mathbf{\Phi}\mathbf{\Sigma}_1\mathbf{\Phi}^T\right) = \mathrm{rank}\left(\mathbf{Q}\right) = n_{\mathbf{\Sigma}_2}$. We can similarly establish that $r_2 = \mathrm{rank}\left(\mathbf{\Phi}\mathbf{\Sigma}_2\mathbf{\Phi}^T\right) = n_{\mathbf{\Sigma}_1}$ and $r_{12} = \mathrm{rank}\left(\mathbf{\Phi}(\mathbf{\Sigma}_i + \mathbf{\Sigma}_j)\mathbf{\Phi}^T\right) = n_{\mathbf{\Sigma}_1} + n_{\mathbf{\Sigma}_2}$, that is, this matrix construction, which satisfies the condition in (72), achieves the maximum diversity-order in (69).

B. Case Where $M < NO_{Dim}$

The maximum diversity-order that we can achieve with any measurement matrix when $M < NO_{Dim}$ is now given by:

$$d \leq \frac{1}{4}M. \tag{76}$$

This upper bound follows from the solution to the following integer-valued optimization problem[10]:

$$\max_{r_1, r_2, r_{12}} -\frac{1}{2}\left(\frac{r_1 + r_2}{2} - r_{12}\right) \tag{77}$$

[10]Note that this problem represents a relaxation of the actual diversity-order maximization problem, as it incorporates only some of the constraints dictated by the geometrical description of the scenario. For example, it does not take into account the actual value of some parameters of the input description as $r_{\mathbf{\Sigma}_1}, r_{\mathbf{\Sigma}_2}$ and $r_{\mathbf{\Sigma}_{12}}$.

subject to: $r_1 + r_2 \geq r_{12}$, $r_1 \leq M$, $r_2 \leq M$, $r_{12} \leq M$ and $r_1, r_2, r_{12} \in \mathbb{Z}_0^+$.

The solution, which can be obtained by considering a linear programming relaxation along with a Branch and Bound approach [84], is given by[11]:

$$r_1 + r_2 = r_{12} \quad , \quad r_{12} = M \tag{78}$$

and

$$d = -\frac{1}{2}\left(\frac{r_1 + r_2}{2} - r_{12}\right) = \frac{1}{4}M. \tag{79}$$

A possible measurement matrix construction that achieves such maximum diversity-order in (79) is obtained by picking arbitrarily only M among its $n_{\mathbf{\Sigma}_1} + n_{\mathbf{\Sigma}_2}$ row vectors of the matrix $\mathbf{\Phi}$ in (73). In particular, we take M_1 rows from the set $\left[\mathbf{v}_1, \ldots, \mathbf{v}_{n_{\mathbf{\Sigma}_1}}\right]$ and M_2 rows from the set $\left[\mathbf{w}_1, \ldots, \mathbf{w}_{n_{\mathbf{\Sigma}_2}}\right]$, where $M_1 + M_2 = M$, which is always possible as $M < NO_{Dim} = n_{\mathbf{\Sigma}_1} + n_{\mathbf{\Sigma}_2}$. Then, by following steps similar to the previous ones, it is possible to show that $r_1 = \operatorname{rank}\left(\mathbf{\Phi}\mathbf{\Sigma}_1\mathbf{\Phi}^T\right) = M_2$, $r_2 = \operatorname{rank}\left(\mathbf{\Phi}\mathbf{\Sigma}_2\mathbf{\Phi}^T\right) = M_1$ and $r_{12} = \operatorname{rank}\left(\mathbf{\Phi}(\mathbf{\Sigma}_i + \mathbf{\Sigma}_j)\mathbf{\Phi}^T\right) = M_1 + M_2 = M$, that is, this matrix construction achieves the maximum diversity-order in (79).

APPENDIX E
PROOF OF THEOREM 5

As presented in Appendix C, the upper bound to the misclassification probability decays exponentially with $1/\sigma^2$, as $\sigma^2 \to 0$ (achieving a diversity-order equal to infinity) if

$$\sum_{i=r_{12}+1}^{M} \mathbf{u}_{12,i}^T \mathbf{\Phi}(\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2)(\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2)^T \mathbf{\Phi}^T \mathbf{u}_{12,i} > 0. \tag{80}$$

However, we underline that, in view of the fact that the condition $M \leq r_{\mathbf{\Sigma}_{12}}$ no longer implies $r_{12} = M$, the number of measurements required to achieve infinite diversity-order with a optimized $\mathbf{\Phi}$ can be lower than those with a random $\mathbf{\Phi}$.

Assume that

$$(\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2) \notin \operatorname{im}(\mathbf{\Sigma}_1 + \mathbf{\Sigma}_2). \tag{81}$$

We can show that that it is possible to achieve infinite diversity-order with the measurement kernel

$$\mathbf{\Phi} = \left[\boldsymbol{\phi}^T\right] \tag{82}$$

where $\boldsymbol{\phi} \in \operatorname{Null}(\mathbf{\Sigma}_1 + \mathbf{\Sigma}_2)$ – note that there exists such a $\boldsymbol{\phi}$ because $\operatorname{Null}(\mathbf{\Sigma}_1 + \mathbf{\Sigma}_2)$ is the orthogonal complement of $\operatorname{im}(\mathbf{\Sigma}_1 + \mathbf{\Sigma}_2)$ and hence it does not contain only the zero vector. In fact, let us consider the standard decomposition [85]

$$(\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2) = \boldsymbol{\mu}_{im} + \boldsymbol{\mu}_{Null}, \tag{83}$$

where $\boldsymbol{\mu}_{im} \in \operatorname{im}(\mathbf{\Sigma}_1 + \mathbf{\Sigma}_2)$, $\boldsymbol{\mu}_{Null} \in \operatorname{Null}(\mathbf{\Sigma}_1 + \mathbf{\Sigma}_2)$ and, given (81), $\boldsymbol{\mu}_{Null} \neq \mathbf{0}$. Then, $\mathbf{\Phi}(\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2) = 0 + \boldsymbol{\phi}^T\boldsymbol{\mu}_{Null} > 0$ and, therefore, $\sum_{i=r_{12}+1}^{M} \mathbf{u}_{12,i}^T \mathbf{\Phi}(\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2)(\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2)^T \mathbf{\Phi}^T \mathbf{u}_{12,i} = |\boldsymbol{\phi}^T\boldsymbol{\mu}_{Null}|^2 > 0$.

[11]Note that the solution of the optimization problem is not unique. Nevertheless, the maximum value achieved by the objective function is indeed unique.

Assume now that

$$(\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2) \in \mathrm{im}\,(\boldsymbol{\Sigma}_1 + \boldsymbol{\Sigma}_2)\,. \tag{84}$$

We can now show that one cannot achieve infinite diversity-order for all possible choices of the measurement kernel because $\boldsymbol{\Phi}(\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2) \in \mathrm{im}\left(\boldsymbol{\Phi}(\boldsymbol{\Sigma}_1 + \boldsymbol{\Sigma}_2)\boldsymbol{\Phi}^T\right)$ and hence $\sum_{i=r_{12}+1}^{M} \mathbf{u}_{12,i}^T \boldsymbol{\Phi}(\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2)(\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2)^T \boldsymbol{\Phi}^T \mathbf{u}_{12,i} = 0$. In fact, in view of this last result, the maximum diversity-order and the measurement kernel that achieves such a diversity-order are now given by Theorem 4.

REFERENCES

[1] E. Candès, J. K. Romberg, and T. Tao, "Stable signal recovery for incomplete and inaccurate measurements," *Pure and Applied Mathematics, Communications on*, vol. 59, no. 8, pp. 1207 –1223, August 2006.

[2] E. Candès, J. Romberg, and T. Tao, "Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information," *Information Theory, IEEE Transactions on*, vol. 52, no. 2, pp. 489–509, 2006.

[3] E. Candès and T. Tao, "Near-optimal signal recovery from random projections: Universal encoding strategies?" *Information Theory, IEEE Transactions on*, vol. 52, no. 12, pp. 5406–5425, 2006.

[4] D. Donoho, "Compressed sensing," *Information Theory, IEEE Transactions on*, vol. 52, no. 4, pp. 1289 –1306, april 2006.

[5] R. Baraniuk, M. Davenport, R. DeVore, and M. Wakin, "A simple proof of the restricted isometry property for random matrices," *Constructive Approximation*, vol. 28, no. 3, pp. 253 –263, December 2008.

[6] E. Candès and T. Tao, "Decoding by linear programming," *Information Theory, IEEE Transactions on*, vol. 51, no. 12, pp. 4203 –4215, december 2005.

[7] S. Mallat and Z. Zhang, "Matching pursuits with time-frequency dictionaries," *Signal Processing, IEEE Transactions on*, vol. 41, no. 12, pp. 3397–3415, 1993.

[8] S. S. Chen, D. L. Donoho, Michael, and A. Saunders, "Atomic decomposition by basis pursuit," *Scientific Computing, SIAM Journal on*, vol. 20, pp. 33–61, 1998.

[9] J. Tropp and S. Wright, "Computational methods for sparse solution of linear inverse problems," *Proceedings of the IEEE*, vol. 98, no. 6, pp. 948–958, 2010.

[10] J. Romberg, "Imaging via compressive sampling," *Signal Processing Magazine, IEEE*, vol. 25, no. 2, pp. 14–20, 2008.

[11] D. Healy and D. Brady, "Compression at the physical interface," *Signal Processing Magazine, IEEE*, vol. 25, no. 2, pp. 67–71, 2008.

[12] M. Lustig, D. Donoho, J. Santos, and J. Pauly, "Compressed sensing MRI," *Signal Processing Magazine, IEEE*, vol. 25, no. 2, pp. 72–82, 2008.

[13] J. Haupt, W. Bajwa, M. Rabbat, and R. Nowak, "Compressed sensing for networked data," *Signal Processing Magazine, IEEE*, vol. 25, no. 2, pp. 92–101, 2008.

[14] M. Davenport, P. Boufounos, M. Wakin, and R. Baraniuk, "Signal processing with compressive measurements," *Selected Topics in Signal Processing, IEEE Journal of*, vol. 4, no. 2, pp. 445–460, 2010.

[15] A. Ashok, P. Baheti, and M. Neifeld, "Compressive imaging system design using task-specific information," *Applied Optics*, vol. 47, no. 25, pp. 4457 –4471, September 2008.

[16] P. K. Baheti and M. A. Neifeld, "Adaptive feature-specific imaging: A face recognition example," *Applied Optics*, vol. 47, no. 10, pp. B21 –B31, April 2008.

[17] P. Baheti and M. Neifeld, "Recognition using information-optimal adaptive feature-specific imaging," *the Optical Society of America A, Journal of*, vol. 26, no. 4, pp. 1055 –1070, 2009.

[18] J. Mairal, F. Bach, J. Ponce, G. Sapiro, and A. Zisserman, "Discriminative learned dictionaries for local image analysis," in *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*, june 2008.

[19] Y. Hitomi, J. Gu, M. Gupta, T. Mitsunaga, and S. Nayar, "Video from a single coded exposure photograph using a learned over-complete dictionary," in *Computer Vision (ICCV), 2011 IEEE International Conference on*, 2011, pp. 287–294.

[20] J. Wright, A. Y. Yang, A. Ganesh, S. Sankar Sastry, and Y. Ma, "Robust face recognition via sparse representation," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 31, no. 2, pp. 210 –227, february 2009.

[21] R. Calderbank, S. Jafarpour, and R. Schapire, "Compressed learning: Universal sparse dimensionality reduction and learning in the measurement domain," Tech. Rep., 2009.

[22] M. Chen, W. Carson, M. R. D. Rodrigues, R. Calderbank, and L. Carin, "Communication inspired linear discriminant analysis," in *Proceedings of the 29th International Conference on Machine Learning*, june 2012.

[23] R. O. Duda, P. E. Hart, and D. G. Stork, *Pattern Classification (2nd Edition)*. New York, NY: Wiley-Interscience, 2000.

[24] R. Fisher, "The use of multiple measurements in taxonomic problems," *Human Genetics, Annals of*, vol. 7, no. 2, pp. 179 –188, 1936.

[25] O. Hamsici and A. Martinez, "Bayes optimality in linear discriminant analysis," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 30, no. 4, pp. 647–657, 2008.

[26] W. R. Carson, M. R. D. Rodrigues, M. Chen, L. Carin, and R. Calderbank, "Communications-inspired projection design with application to compressive sensing," *Imaging Sciences, SIAM Journal on*, vol. 5, no. 4, pp. 1185 –1212, 2012.

[27] D. Erdogmus and J. C. Principe, "Lower and upper bounds for misclassification probability based on Renyi's information," *VLSI signal processing systems for signal, image and video technology, Journal of*, vol. 37, no. 2 -3, pp. 305–317, 2004.

[28] K. Hild, D. Erdogmus, K. Torkkola, and J. Principe, "Feature extraction using information-theoretic learning," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 28, no. 9, pp. 1385–1392, 2006.

[29] S. Kaski and J. Peltonen, "Informative discriminant analysis," in *Proceedings of the Twentieth International Conference on Machine Learning (ICML)*, 2003, pp. 329–336.

[30] L. Liu and P. Fieguth, "Texture classification from random features," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 34, no. 3, pp. 574 –586, march 2012.

[31] Z. Nenadic, "Information discriminant analysis: Feature extraction with an information-theoretic objective," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 29, no. 8, pp. 1394–1407, august 2007.

[32] D. Tao, X. Li, X. Wu, and S. Maybank, "Geometric mean for subspace selection," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 31, no. 2, pp. 260–274, 2009.

[33] K. Torkkola, "Learning discriminative feature transforms to low dimensions in low dimensions," in *Advances in neural information processing systems 14 (NIPS)*. MIT Press, 2001, pp. 3–8.

[34] ——, "Feature extraction by non-parametric mutual information maximization," *Machine Learning Research, Journal of*, vol. 3, pp. 1415–1438, March 2003.

[35] J. Tenenbaum, V. de Silva, and J. Langford, "A global geometric framework for nonlinear dimensionality reduction," *Science*, vol. 290, no. 5500, pp. 2319–2323, December 2000.

[36] R. Calderbank and S. Jafarpour, "Finding needles in compressed haystacks," in *Proceedings of 2012 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, March 2012, pp. 3441–3444.

[37] M. Duarte, M. Davenport, M. Wakin, J. Laska, D. Takhar, K. Kelly, and R. Baraniuk, "Multiscale random projections for compressive classification," in *Proceedings of the International Conference on Image Processing*, vol. 6, 2007.

[38] M. A. Davenport, R. M. F. Duarte, R. M. B. Wakin, C. J. N. L. R, D. Takhar, R. K. F. Kelly, and R. R. G. B. R, "The smashed filter for compressive classification and target recognition," in *Proceedings of SPIE Symposium on Electronic Imaging: Computational Imaging*, 2007, p. 6498.

[39] M. Duarte, M. Davenport, M. Wakin, and R. Baraniuk, "Sparse signal detection from incoherent projections," in *Proceedings of 2006 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, vol. 3, 2006.

[40] J. Haupt, R. Castro, R. Nowak, G. Fudge, and A. Yeh, "Compressive sampling for signal classification," in *Proceedings of Asilomar Conference on Signals, Systems and Computers*, 2006, pp. 1430–1434.

[41] J. Haupt and R. Nowak, "Compressive sampling for signal detection," in *Proceedings of 2007 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, vol. 3, 2007.

[42] T. Wimalajeewa, H. Chen, and P. Varshney, "Performance limits of compressive sensing-based signal classification," *Signal Processing, IEEE Transactions on*, vol. 60, no. 6, pp. 2758–2770, June 2012.

[43] K. Krishnamurthy, M. Raginsky, and R. Willett, "Hyperspectral target detection from incoherent projections," in *Proceedings of 2010 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2010, pp. 3550–3553.

[44] ——, "Hyperspectral target detection from incoherent projections: Nonequiprobable targets and inhomogeneous snr," in *Proceedings of 2010 IEEE International Conference on Image Processing (ICIP)*, 2010, pp. 1357–1360.

[45] M. A. Neifeld, A. Ashok, and P. K. Baheti, "Task-specific information for imaging system analysis," *Optical Society of America A, Journal of the*, vol. 24, no. 12, pp. B25–B41, December 2007.

[46] J. Ke, A. Ashok, and M. A. Neifeld, "Object reconstruction from adaptive compressive measurements in feature-specific imaging," *Applied Optics*, vol. 49, no. 34, pp. H27–H39, Dec 2010.

[47] J. M. Duarte-Carvajalino, G. Sapiro, G. Yu, and L. Carin, "Online adaptive statistical compressed sensing of Gaussian mixture models," *CoRR*, vol. abs/1112.5895, 2011.

[48] J. M. Duarte-Carvajalino, G. Yu, L. Carin, and G. Sapiro, "Task-driven adaptive statistical compressive sensing of Gaussian mixture models," *Signal Processing, IEEE Transactions on*, vol. 61, no. 3, pp. 585–600, 2013.

[49] S. Aeron, V. Saligrama, and M. Zhao, "Information theoretic bounds for compressed sensing," *Information Theory, IEEE Transactions on*, vol. 56, no. 10, pp. 5111–5130, 2010.

[50] M. Akcakaya and V. Tarokh, "Shannon-theoretic limits on noisy compressive sampling," *Information Theory, IEEE Transactions on*, vol. 56, no. 1, pp. 492–504, 2010.

[51] A. Fletcher, S. Rangan, and V. Goyal, "Necessary and sufficient conditions for sparsity pattern recovery," *Information Theory, IEEE Transactions on*, vol. 55, no. 12, pp. 5758–5772, 2009.

[52] K. Rad, "Nearly sharp sufficient conditions on exact sparsity pattern recovery," *Information Theory, IEEE Transactions on*, vol. 57, no. 7, pp. 4672–4679, 2011.

[53] G. Reeves, "Sparsity pattern recovery in compressed sensing," Ph.D. dissertation, Department of Electrical Engineering and Computer Science, Univ. California, Berkeley, 2011.

[54] G. Reeves and M. Gastpar, "The sampling rate-distortion tradeoff for sparsity pattern recovery in compressed sensing," *Information Theory, IEEE Transactions on*, vol. 58, no. 5, pp. 3065–3092, 2012.

[55] ——, "Approximate sparsity pattern recovery: Information-theoretic lower bounds," *Information Theory, IEEE Transactions on*, vol. 59, no. 6, pp. 3451–3465, 2013.

[56] G. Tang and A. Nehorai, "Performance analysis for sparse support recovery," *Information Theory, IEEE Transactions on*, vol. 56, no. 3, pp. 1383–1399, 2010.

[57] A. Tulino, G. Caire, S. Verdú, and S. Shamai, "Support recovery with sparsely sampled free random matrices," *Information Theory, IEEE Transactions on*, vol. 59, no. 7, pp. 4243–4271, 2013.

[58] W. Wang, M. Wainwright, and K. Ramchandran, "Information-theoretic limits on sparse signal recovery: Dense versus sparse measurement matrices," *Information Theory, IEEE Transactions on*, vol. 56, no. 6, pp. 2967–2979, 2010.

[59] M. Wainwright, "Sharp thresholds for high-dimensional and noisy sparsity recovery using l_1-constrained quadratic programming (lasso)," *Information Theory, IEEE Transactions on*, vol. 55, no. 5, pp. 2183–2202, 2009.

[60] ——, "Information-theoretic limits on sparsity recovery in the high-dimensional and noisy setting," *Information Theory, IEEE Transactions on*, vol. 55, no. 12, pp. 5728–5741, 2009.

[61] P. Zhao and B. Yu, "On model selection consistency of lasso," *Machine Learning Research, Journal of*, vol. 7, pp. 2541–2563, December 2006.

[62] V. Tarokh, N. Seshadri, and R. Calderbank, "Space-time codes for high data rate wireless communication: performance criterion and code construction," *Information Theory, IEEE Transactions on*, vol. 44, no. 2, pp. 744 –765, march 1998.

[63] V. Tarokh, H. Jafarkhani, and R. Calderbank, "Space-time block codes from orthogonal designs," *Information Theory, IEEE Transactions on*, vol. 45, no. 5, pp. 1456–1467, 1999.

[64] D. Tse and P. Viswanath, *Fundamentals of Wireless Communication*. Cambridge, U.K.: Cambridge University Press, 2005.

[65] L. Zheng and D. Tse, "Diversity and multiplexing: a fundamental tradeoff in multiple-antenna channels," *Information Theory, IEEE Transactions on*, vol. 49, no. 5, pp. 1073 –1096, may 2003.

[66] M. Nokleby, R. Calderbank, and M. R. D. Rodrigues, "Information-theoretic limits on the classification of gaussian mixtures: Classification on the Grassmann manifold," in *Proceedings of IEEE Information Theory Workshop*, September 2013.

[67] L. Zheng and D. Tse, "The diversity-multiplexing tradeoff for non-coherent multiple antenna channels," in *Proceedings of the Annual Allerton Conference on Communication, Control and Computing*, 2002.

[68] M. Chen, J. Silva, J. Paisley, C. Wang, D. Dunson, and L. Carin, "Compressive sensing on manifolds using a nonparametric mixture of

factor analyzers: Algorithm and performance bounds," *Signal Processing, IEEE Transactions on*, vol. 58, no. 12, pp. 6140 –6155, december 2010.

[69] G. Yu and G. Sapiro, "Statistical compressed sensing of Gaussian mixture models," *Signal Processing, IEEE Transactions on*, vol. 59, no. 12, pp. 5842–5858, 2011.

[70] G. Yu, G. Sapiro, and S. Mallat, "Solving inverse problems with piecewise linear estimators: From Gaussian mixture models to structured sparsity," *Image Processing, IEEE Transactions on*, vol. 21, no. 5, pp. 2481–2499, 2012.

[71] S. Ji, Y. Xue, and L. Carin, "Bayesian compressive sensing," *Signal Processing, IEEE Transactions on*, vol. 56, no. 6, pp. 2346–2356, 2008.

[72] T. Blumensath and M. Davies, "Sampling theorems for signals from the union of finite-dimensional linear subspaces," *Information Theory, IEEE Transactions on*, vol. 55, no. 4, pp. 1872–1882, 2009.

[73] M. Stojnic, F. Parvaresh, and B. Hassibi, "On the reconstruction of block-sparse signals with an optimal number of measurements," *Signal Processing, IEEE Transactions on*, vol. 57, no. 8, pp. 3075–3085, 2009.

[74] Y. Eldar and M. Mishali, "Robust recovery of signals from a structured union of subspaces," *Information Theory, IEEE Transactions on*, vol. 55, no. 11, pp. 5302–5316, 2009.

[75] Y. Eldar, P. Kuppinger, and H. Bolcskei, "Block-sparse signals: Uncertainty relations and efficient recovery," *Signal Processing, IEEE Transactions on*, vol. 58, no. 6, pp. 3042–3054, 2010.

[76] R. G. Baraniuk, V. Cevher, M. F. Duarte, and C. Hegde, "Model-based compressive sensing," *IEEE Trans. Inf. Theory*, vol. 56, no. 4, pp. 1982–2001, Apr. 2010.

[77] R. G. Baraniuk and M. Wakin, "Random projections of smooth manifolds," *Found. Comput. Math.*, vol. 9, no. 1, pp. 51–77, 2009.

[78] Y. Wu and S. Verdú, "Optimal phase transitions in compressed sensing," *Information Theory, IEEE Transactions on*, vol. 58, no. 10, pp. 6241–6263, 2012.

[79] F. Renna, R. Calderbank, L. Carin, and M. R. D. Rodrigues, "Reconstruction of signals drawn from a Gaussian mixture from noisy compressive measurements: MMSE phase transitions and beyond," *arXiv preprint arXiv:1307.0861*, 2013.

[80] A. Paulraj, R. Nabar, and D. Gore, *Introduction to Space-Time Wireless Communications*. Cambridge, U.K.: Cambridge University Press, 2008.

[81] A. Bhattacharyya, "On a measure of divergence between two statistical populations defined by their probability distributions," *Calcutta Mathematical Society, Bulletin of the*, vol. 45, pp. 99 –109, 1943.

[82] Q. Qiu and G. Sapiro, "Learning transformations for clustering and classification," *CoRR*, vol. abs/1309.2074, 2013.

[83] G. H. Golub and C. F. Van Loan, *Matrix Computations (3rd Edition)*. Baltimore, MD: Johns Hopkins University Press, 1996.

[84] A. Schrijver, *Theory of linear and integer programming*. West Sussex, UK: John Wiley and Sons, 1998.

[85] R. Horn and C. Johnson, *Matrix Analysis*. Cambridge, UK: Cambridge University Press, 1985.